ABSTRACT

Title of dissertation: ENUMERATION OF HARMONIC FRAMES AND FRAME BASED DIMENSION REDUCTION

Matthew J. Hirn, Doctor of Philosophy, 2009

Dissertation directed by: Professor John J. Benedetto
Professor Kasso A. Okoudjou
Department of Mathematics

We investigate two aspects of frame theory, one of a theoretical nature, the other very much on the applied side. In the former, we enumerate all harmonic frames of prime order, and develop partial proofs concerning the structure of the symmetry group for this subset of frames. In the latter, we develop frame theory in the context of kernel eigenmap methods, merging the two theories in a practical manner and applying new algorithms to hyperspectral imagery data for the purposes of material classification. These two problems, while seemingly separate, are united by frame theory and serve to illustrate both the beautiful theoretical nature of frames as well as their practicality in dealing with real world problems.

ENUMERATION OF HARMONIC FRAMES AND FRAME BASED DIMENSION REDUCTION

by

Matthew J. Hirn

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2009

Advisory Committee:
Professor John J. Benedetto, Chair/Advisor
Professor Kasso A. Okoudjou, Co-Chair/Co-Advisor
Professor Lawrence C. Washington
Professor Radu V. Balan
Professor Rama Chellappa

Dedication

For my parents, Darleen and John Hirn.

Acknowledgments

This document is the culmination of five years of work, almost all of which was collaborative in some way. I hope here to express in some small way, my enduring gratitude to those that made this extraordinary journey possible.

I must start with my advisor, Professor John Benedetto. John has been, and continues to be, more than I could have ever hoped for in an advisor. When I first came to the University of Maryland, one of the first people I met was John. Since that day he has guided me through the foundations of harmonic analysis and transitioned me into research. John has inspired me with countless ideas and problems, and has always held me accountable to the details. He has pushed me to be the best that I can be, and for that I will always be grateful. Above and beyond what he has passed onto me mathematically though, I have always felt that John has kept my best interests at heart, and that he is always looking out for me. I would not be where I am today without his guidance.

My co-advisor, Kasso Okoudjou, has also been instrumental in my success. I first met Kasso as an undergraduate at Cornell University - I was a senior and he was in the first year of an assistant professor position. At the time I was writing my senior thesis - the first math paper I had ever attempted to write. Kasso took the time to read through my paper, give me suggestions, and also force me to really explain things in clear detail. A few years later when I found out he was coming to the University of Maryland, I knew that I wanted him as my co-advisor. Whether it has been harmonic frames or compressed sensing, Kasso has always been there as someone to help me work out the tough spots and bounce ideas off of. I have always enjoyed just stopping by his office to see if he is there, whether it is to ask a serious math question or simply say hi and see how things are going. And though I have now finished my thesis, I look forward to our continued friendship, and of course working on our next idea.

I would of course like to thank the remaining three members of my committee for generously sparing their time. Professor Radu Balan has always challenged and inspired me with his insightful remarks and questions. Furthermore, I would also just like to say that his course on financial mathematics was one of the most interesting and engaging classes during my time at Maryland. Professor Washington, not knowing me very well, took the time to read an early draft of chapter 2 on harmonic frames. And Professor Chellappa, not knowing me at all, agreed to serve as my dean's representative. I am humbled by their kindness.

I must also thank Professor Wojciech Czaja. Wojtek, as he likes to be called, has been my unofficial third advisor. Mathematically speaking, in the course of our work with hyperspectral and high dimensional data, Wojtek has always challenged me to explore a multitude of different ideas, many with great success. It has been a pleasure to work for and with him. Aside from mathematics though, Wojtek has truly been a good friend. Some of the first conferences I attended were with Wojtek, and in particular I will always remember our Los Angeles trips fondly. Wojtek has generously invited me over during holidays when he knew I had nowhere else to go, and has been a source of encouragement for me during some of the more difficult times in the past year.

Most importantly I would like to give a heartfelt thank you to my parents, Darleen and John Hirn. Without them none of this would have been possible. Their love and support throughout these five years has been incredible. They have made many sacrifices to put me through my undergraduate education, and even through graduate school have still partially supported me, for which I cannot thank them enough. It has been their continued belief in me though, that has truly kept me going even in the hard times.

My new wife Yuan Yuan Sun has been the person I have turned to the most. Through all of the ups and downs, Yuan Yuan has been there for me, constantly supporting me through this endeavor. Whenever I would discover something new or get excited about an idea, I could not wait to tell Yuan Yuan; conversely, when I was feeling down or as if things were too tough, she was always there to pick me up. It has been a great joy for me to be able to share this accomplishment with her, and I cannot wait to see where our paths will take us next.

To my brother Bradley Hirn, I am forever grateful for your unwavering support. Our conversations would always leave me in higher spirits, and I knew that you would always be there for me.

Of course where would one be without his friends. First and foremost, I would like to acknowledge my officemates for the better part of these last five years, Lucas Vaczlavik and Stefan Mendez-Diez. Lucas and Stefan have been true friends these last five years, and I feel lucky and honored to call them my friends. I will forever look fondly upon our time here at Maryland together, whether it be sitting around the office distracting each other or hanging out on the weekends. I cannot wait to see where life takes us next. I would also like mention John Habert and Toni Watson, both of who's friendship I value to this day. A significant portion of my work has been with my fellow students and postdocs on the NGA project - Martin Ehler, Chris Flake, Nate Strawn, Jesse Sugar-Moore, and David Widemann - it has been a pleasure to work with all of them. Finally I would like mention all of my office mates in 4423 from the first three years - it has been fun watching all of us progress through the program, and now we have made it.

There are many other people at Maryland that have helped me along the way. Ioannis Konstantinidis for his help with any and all math problems - he always seemed to know the answer; William Schildknecht for his help and guidance when I was a teaching assistant; Fletcher Kinne for his general helpfulness with the copier or the fax machine and a host of other things; the ladies of the graduate office - Linette Berry, Haydee Hidalgo, and Celeste Regalado - who always put a smile on my face; the staff in the business office - Jamie Carrigan, Liliana Gonzalez, Rachel Katz, Sharon Welton, and Elizabeth Wincek - for always being generous with their time and helping me through any issue I may have had, especially during the year in which I helped organize the NWC seminar; and finally the people from the mathnet systems department staff - Mark Tilmes, Tony Zhang, and Taizhu Zhou, who were always able to answer any of my questions.

There are of course countless other people that I have not named who have each in their own way been a part of this journey - I thank all of you.

Lastly I would like to thank God for blessing me with this great opportunity - I truly am one of the lucky ones.

Table of Contents

List of Figures

Chapter 1

Introduction to Frames

1.1 Frames

Given a finite dimensional Hilbert space $\mathbb{H}$, a basis is a set of elements that give a unique representation for each element $\mathbb{H}$. Frames, on the other hand, are an overcomplete set of elements that allow for an infinite number of representations of each element in $\mathbb{H}$. While bases are useful in certain situations because the representation is unique, at other times it is better to have the flexibility provided by a frame.

The purpose of this dissertation is two-fold. In chapter two we examine a subclass of finite frames known as harmonic frames. In particular, we will study harmonic frames with a prime number of elements. The main result is to prove a recursive formula for the number of harmonic frames of prime order. A secondary result is to partially determine the symmetry group of all such harmonic frames.

The second main focus is to examine the usefullness of frames, in conjunction with kernel based dimension reducing methods, for the classification of materials in multispectral and hyperspectral imagery data. Results here are theoretically motivated, yet are empirical in nature. We plan to give results that exhibit the promise of this approach. The theoretical motivations can be found in chapter three, while chapter four contains empirical results.

First though, we begin with an introduction to frame theory.

1.2 Frame Theory

Let $\mathcal{I}$ be a possibly infinite, but countable, index set. A *frame* [24, 25, 14, 22] for a separable Hilbert space $\mathbb{H}$ is a collection of vectors

$$\Phi = \{\varphi_i : i \in \mathcal{I}\} \subset \mathbb{H} \tag{1.1}$$

for which there exists constants $0 < A \leq B < \infty$ such that for each $f \in \mathbb{H}$,

$$A\|f\|^2 \leq \sum_{i \in \mathcal{I}} |\langle \varphi_i, f \rangle|^2 \leq B\|f\|^2. \tag{1.2}$$

Constants A and B which satisfy (1.2) are called *frame bounds* of Φ. Optimally chosen values of A and B are referred to as the *optimal frame bounds* of the frame. When $A = B$, the frame Φ is referred to as a *tight frame.*

As an example of a frame one may choose an orthonormal basis - it is in fact a tight frame with constants $A = B = 1$. A union of any two orthonormal bases is a tight frame with constants $A = B = 2$, etc. A union of an orthonormal basis with N arbitrary unit norm vectors is a frame with bounds $A = 1$ and $B = N + 1$. If the Hilbert space is infinite dimensional and N is finite this last example is certainly not a tight frame. Some other examples are given by figures 1.1, 1.2, and 1.3.

Given a frame $\Phi = \{\varphi_i : i \in \mathcal{I}\}$, a *dual frame* is a collection of vectors $\widehat{\Phi} = \{\hat{\varphi}_i : i \in \mathcal{I}\} \subset \mathbb{H}$ such that for all $f \in \mathbb{H}$, we have the reconstruction formula

$$f = \sum_{i \in \mathcal{I}} \langle f, \varphi_i \rangle \hat{\varphi}_i. \tag{1.3}$$

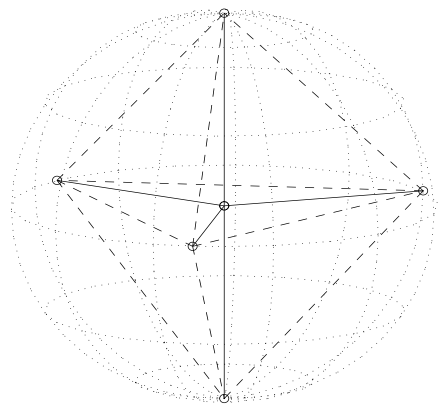

Figure 1.1: Frame with six elements in $\mathbb{R}^3$

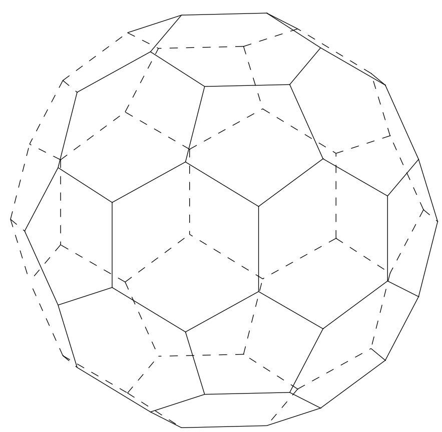

Figure 1.2: Buckyball tight frame

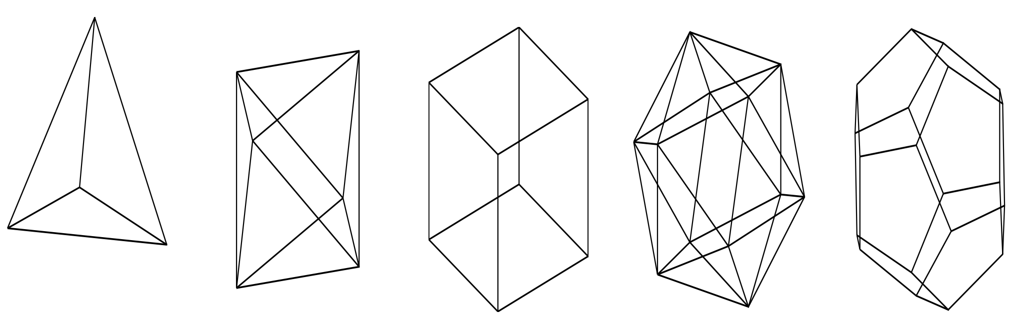

Figure 1.3: The platonic solids form tight frames

It is perhaps not immediately clear that every frame should have a dual frame. In order to obtain a dual frame to a frame Φ, we will define the frame operator. Let $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$, and define $\ell^2(\mathcal{I})$ as the space of all sequences indexed by $\mathcal{I}$ with finite energy, i.e.

$$\ell^2(\mathcal{I}) := \{c = (c_i)_{i\in\mathcal{I}} : c_i \in \mathbb{F} \;\; \forall\, i \in \mathcal{I} \text{ and } \sum_{i\in\mathcal{I}} |c_i|^2 < \infty\}. \tag{1.4}$$

Given a frame $\Phi = \{\varphi_i : i \in \mathcal{I}\}$, the *analysis operator* $L : \mathbb{H} \to \ell^2(\mathcal{I})$ is defined by

$$L(f) := \big(\langle f, \varphi_i\rangle\big)_{i\in\mathcal{I}}. \tag{1.5}$$

The adjoint of the analysis operator $L^\star$ is called the *synthesis operator,* and $S = L^\star L$ is the *frame operator.* For each $c \in \ell^2(\mathcal{I})$, the synthesis operator is defined by

$$L^\star(c) = \sum_{i\in\mathcal{I}} c_i\varphi_i. \tag{1.6}$$

Given the previous two equations, it is easy to see that the frame operator is defined by

$$S(f) = \sum_{i\in\mathcal{I}} \langle f, \varphi_i\rangle\varphi_i, \tag{1.7}$$

where $f \in \mathbb{H}$. The following known theorems characterize the analysis, synthesis, and frame operators.

Theorem 1.2.1 ([16])**.** *Let $\Phi = \{\varphi_i : i \in \mathcal{I}\} \subset \mathbb{H}$ be a frame for $\mathbb{H}$. Then the following are satisfied:*

a. *L is a bounded operator from $\mathbb{H}$ into $\ell^2(\mathcal{I})$.*

b. *$L^\star$ extends to a bounded operator from $\ell^2(\mathcal{I})$ into $\mathbb{H}$.*

c. *L and $L^\star$ are adjoint operators of each other.*

Theorem 1.2.2 ([16])**.** *Let $\Phi = \{\varphi_i : i \in \mathcal{I}\} \subset \mathbb{H}$ be a frame for $\mathbb{H}$. The frame operator $S = L^\star L$ maps $\mathbb{H}$ onto $\mathbb{H}$ and is a positive invertible operator satisfying $A \cdot Id \leq S \leq B \cdot Id$ and $B^{-1} \cdot Id \leq S^{-1} \leq A^{-1} \cdot Id$. In particular, Φ is a tight frame if and only if $S = A \cdot Id$.*

Note, in theorem 1.2.2 Id denotes the identity map on $\mathbb{H}$, i.e., $Id(f) = f$ for all $f \in \mathbb{H}$. The sequence of vectors $\{S^{-1}(\varphi_i) : i \in \mathcal{I}\}$ is called the *canonical dual frame*, and is a dual frame for $\Phi = \{\varphi_i : i \in \mathcal{I}\}$. That is we have

$$f = \sum_{i \in \mathcal{I}} \langle f, S^{-1}(\varphi_i) \rangle \varphi_i \tag{1.8}$$

and

$$f = \sum_{i \in \mathcal{I}} \langle f, \varphi_i \rangle S^{-1}(\varphi_i), \tag{1.9}$$

where both sums converge unconditionally in $\mathbb{H}$.

We note here that dual frames are not in general unique and this underlines the importance of the canonical dual frame.

For a particular given frame, it may not be easy to apply the procedure in the preceding paragraph to obtain a dual frame. One special case in which it is easy is that of Parseval frames. A *Parseval frame* is a tight frame consisting of unit norm vectors. If $\Phi = \{\varphi_i : i \in \mathcal{I}\}$ is a Parseval frame, then for every $f \in \mathbb{H}$,

$$f = \sum_{i \in \mathcal{I}} \langle f, \varphi_i \rangle \varphi_i. \tag{1.10}$$

In particular, Parseval frames are dual frames of themselves. For this reason, among others, Parseval frames are the 'best behaved' of frames, and we will present here some of their additional properties.

Most of the basic, general properties of Parseval frames can be derived from the following.

Theorem 1.2.3 ([16]). *A collection of vectors $\Phi = \{\varphi_i : i \in \mathcal{I}\} \subset \mathbb{H}$ is a Parseval frame for $\mathbb{H}$ if and only if there exists a Hilbert space $\mathbb{K}$ containing $\mathbb{H}$ as a closed subspace and an orthonormal basis $\{e_i : i \in \mathcal{I}\}$ of $\mathbb{K}$ such that for all $i \in \mathcal{I}$, $Pe_i = \varphi_i$, where P is the orthogonal projection onto $\mathbb{H}$.*

Equation (1.10) follows immediately from Theorem 1.2.3. Indeed, we have for $f \in \mathbb{H}$,

$$\begin{aligned} P^2 f &= P(Pf) = \sum_{i \in \mathcal{I}} \langle Pf, e_i \rangle Pe_i \\ &= \sum_{i \in \mathcal{I}} \langle f, Pe_i \rangle \varphi_i \\ &= \sum_{i \in \mathcal{I}} \langle f, \varphi_i \rangle \varphi_i. \end{aligned}$$

1.3 Finite Frame Theory

In finite dimensional Hilbert vector spaces, the notion of a frame becomes intutively simple. Let $s, d \in \mathbb{N}$, and suppose $s \geq d$; $\Phi = \{\varphi_i : i = 1, \ldots, s\}$ is a frame for $\mathbb{F}^d$ (recall $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$) if and only if it is a spanning system for $\mathbb{F}^d$. In the finite setting it is often convenient to use matrix notation when working with frames. As such, we will consider φ_j as a vector in $\mathbb{F}^d$, and Φ as a $d \times s$ matrix, where the j^{th} column is φ_j. More explicitly:

$$\varphi_j = \left(\varphi_j(i)\right)_{i=1}^{d} \tag{1.11}$$

$$\Phi \in \mathcal{M}_{d\times s}(\mathbb{F}) \quad \text{and} \quad \Phi_{i,j} = \varphi_j(i). \tag{1.12}$$

Recasting section 1.2 in terms of finite frames and matrices, we see that the analysis operator, L, now maps $\mathbb{F}^d$ into $\mathbb{F}^s$. In fact, for each $f \in \mathbb{F}^d$, the analysis operator is given by:

$$L(f) = \Phi^\star f = \left(\langle f, \varphi_i\rangle\right)_{i=1}^{s}. \tag{1.13}$$

Similarly, the synthesis operator maps $\mathbb{F}^s$ onto $\mathbb{F}^d$, and for each $c \in \mathbb{F}^s$ is given by:

$$L^\star(c) = \Phi c = \sum_{i=1}^{s} c_i\varphi_i. \tag{1.14}$$

Combining equations (1.13) and (1.14), we see that the frame operator maps $\mathbb{F}^d$ to $\mathbb{F}^d$ and, for each $f \in \mathbb{F}^d$, is given by:

$$S(f) = L^\star L(f) = \Phi\Phi^\star f = \sum_{i=1}^{s}\langle f, \varphi_i\rangle\varphi_i. \tag{1.15}$$

A frame that is finite, tight, and unit norm is known as a *finite unit norm tight frame*, or a *FUNTF*. If Φ is a FUNTF with frame constant A, then it is known that $A = s/d$ and $S = \frac{s}{d}I$, where I is the $d \times d$ identity matrix.

One way to characterize FUNTFs is the frame potential [8]. Let $\mathbb{S}^{d-1} \subset \mathbb{F}^d$ denote the unit sphere in $\mathbb{F}^d$. For any unit norm frame $\Phi = \{\varphi_i : i = 1, \ldots, s\}$, the frame potential is defined as

$$FP : \underbrace{\mathbb{S}^{d-1} \times \ldots \times \mathbb{S}^{d-1}}_{s \text{ times}} \rightarrow [0, \infty) \tag{1.16}$$

$$FP(\Phi) := \sum_{i,j=1}^{s} |\langle\varphi_i, \varphi_j\rangle|^2. \tag{1.17}$$

The following theorem characterizes FUNTFs in terms of the frame potential.

Theorem 1.3.1 ([8]). *For a given s and d, the following hold:*

a. *Every local minimizer of the frame potential is also a global minimizer.*

b. *If $s \leq d$, the minimum value of the frame potential is s, and the minimizers are precisely the orthonormal sequences in $\mathbb{C}^d$.*

c. *If $s \geq d$, the minimum value of the frame potential is s^2/d, and the minimizers are precisely the FUNTFs for $\mathbb{C}^d$.*

1.4 Finite Subspace Frames[1]

In the finite frame setting, it is not difficult to define the notion of a finite subspace frame. Let $\Phi = \{\varphi_i : i = 1, \dots, s\} \subset \mathbb{F}^d$ and let W be a subspace of $\mathbb{F}^d$ of dimension $r < d$. We say Φ is a *finite subspace frame* for W if $\operatorname{span}(\Phi) = W$. It is clear from this definition that there exist constants $0 < A \leq B < \infty$ such that for each $f \in W$,

$$A\|f\|^2 \leq \sum_{i=1}^{s} |\langle f, \varphi_i \rangle|^2 \leq B\|f\|^2. \tag{1.18}$$

We note that if we had instead used (1.18) as our definition, then it would not necessarily imply that $\operatorname{span}(\Phi) = W$ but rather that $\operatorname{span}(\Phi) \supseteq W$. The unit norm property as well as the notion of a tight frame remain similar in this setting. More specifically, if we can take $A = B$ in (1.18) then we call Φ a *tight subspace frame*.

[1] While this section is almost certainly not original, it was independently co-authored by David Widemann, University of California at Davis, and the author.

Finally, if Φ is a finite unit norm tight subspace frame, then we say Φ is a *subspace FUNTF.*

We define L, $L^\star$, and S exactly the same as we did previously, however we note that the properties of these maps change for subspace frames. In particular, we see:

(a) $L : \mathbb{F}^d \to \mathbb{F}^s$ is no longer injective, but rather $\ker(L) = (\mathbb{F}^d \setminus W) \cup \{0\}$.

(b) $L^\star : \mathbb{F}^s \to \mathbb{F}^d$ is no longer surjective, but rather $\text{image}(L^\star) = W$.

(c) Based on (a) and (b), we see that $S : \mathbb{F}^d \to \mathbb{F}^d$ is no longer invertible.

Because of (c), theorem 1.2.2 nor equations (1.8) and (1.9) hold for subspace frames. The question then becomes: in what sense do subspace frames behave like standard frames? Theorems below show that subspace frames satisfy natural modifications of theorem 1.2.2, equation (1.8), and equation (1.9).

Let W_{on} be a set of r orthonormal vectors such that $\text{span}(W_{on}) = W$. We will also consider W_{on} as an $d \times r$ matrix where the columns of this matrix are the vectors in the set W_{on}. We define Φ_W to be the $r \times s$ matrix whose columns are the coordinates of Φ in W_{on}; that is:

$$\Phi_W := W_{on}^\star \Phi. \tag{1.19}$$

The i^{th} column of Φ_W is the projected W-subspace coordinates of φ_i.

Proposition 1.4.1. *The columns of Φ_W are a frame for $\mathbb{F}^r$.*

Proof. Since $\text{span}(W_{on}) = W$, we have $\ker(W_{on}^\star) \cap W = \{0\}$. Therefore, since $\text{span}(\Phi) = W$ as well, we see that $W_{on}^\star \Phi$ has rank r. $\square$

We denote the analysis, synthesis, and frame operators of Φ_W by L_W, $L_W^\star$, and S_W, respectively. In terms of the analysis operator, L, for Φ, we have for each $g \in \mathbb{F}^r$,

$$L_W(g) = L(W_{on}g) = \Phi^\star W_{on}g. \tag{1.20}$$

Similarly, for each $c \in \mathbb{F}^s$, the synthesis operator of Φ_W is defined as

$$L_W^\star(c) = W_{on}^\star L^\star(c) = W_{on}^\star \Phi c. \tag{1.21}$$

Combining equations (1.20) and (1.21) we see that for each $g \in \mathbb{F}^r$, S_W is defined as

$$S_W(g) = L_W^\star L_W(g) = W_{on}^\star L^\star(L(W_{on}g)) = W_{on}^\star \Phi\Phi^\star W_{on}g = W_{on}^\star S W_{on}g. \tag{1.22}$$

By proposition 1.4.1 we see that S_W will satisfy theorem 1.2.2 as well as equations (1.8) and (1.9).

Theorem 1.4.2. *Φ is a subspace FUNTF for W with frame bound A if and only if Φ_W is a FUNTF for $\mathbb{F}^r$ with frame bound A.*

Proof. We do the forward direction first: let $g \in \mathbb{F}^r$, then:

$$\begin{aligned}
\langle S_W g, g\rangle &= \langle L_W g, L_W g\rangle \\
&= \langle \Phi^\star W_{on}g, \Phi^\star W_{on}g\rangle \\
&= \sum_{j=1}^{s} |\langle W_{on}g, \varphi_j\rangle|^2 \\
&= A\|W_{on}g\|^2 \\
&= A\langle W_{on}g, W_{on}g\rangle
\end{aligned}$$

Therefore we have:

$$\begin{aligned}
\langle S_W g, g\rangle - A\langle W_{on}g, W_{on}g\rangle &= 0 \quad \Longrightarrow \\
\langle S_W g, g\rangle - A\langle W_{on}^{\star} W_{on}g, g\rangle &= 0 \quad \Longrightarrow \\
\langle g, (S_W - AI)g\rangle &= 0 \quad \Longrightarrow \\
S_W &= AI.
\end{aligned}$$

For the reverse direction, let $f \in W$. There exists $g \in \mathbb{F}^r$ such that $W_{on}g = f$. Therefore,

$$\begin{aligned}
A\|f\|^2 &= A\langle f, f\rangle \\
&= A\langle W_{on}g, W_{on}g\rangle \\
&= \langle Ag, g\rangle \\
&= \langle S_W g, g\rangle \\
&= \langle W_{on}^{\star}\Phi\Phi^{\star}W_{on}g, g\rangle \\
&= \langle \Phi^{\star}W_{on}g, \Phi^{\star}W_{on}g\rangle \\
&= \langle \Phi^{\star}f, \Phi^{\star}f\rangle \\
&= \sum_{j=1}^{s} |\langle f, \varphi_j\rangle|^2
\end{aligned}$$

$\square$

Corollary 1.4.3. *If Φ is a subspace FUNTF for W with frame bound A, then $A = s/r$.*

We define the canonical dual frame of Φ_W in the usual way, that is $\widehat{\Phi}_W =$

$S_W^{-1}\Phi_W$. We now define the *canonical dual subspace frame* of Φ as follows:

$$\widehat{\Phi} = W_{on}\widehat{\Phi}_W = W_{on}S_W^{-1}W_{on}^\star\Phi. \tag{1.23}$$

As the name implies, the set $\widehat{\Phi} = \{\hat{\varphi}_i : i = 1, \dots s\} = \{W_{on}S_W^{-1}W_{on}^\star\varphi_i : i = 1, \dots, s\}$ will have the following properties:

Proposition 1.4.4. *$\widehat{\Phi}$ is a subspace frame for W.*

Proof. This follows from proposition 1.4.1. □

Theorem 1.4.5. *Every $f \in \mathrm{W}$ can be represented as*

$$f = \sum_{i=1}^{s}\langle f, \hat{\varphi}_i\rangle\varphi_i = \sum_{i=1}^{s}\langle f, \varphi_i\rangle\hat{\varphi}_i. \tag{1.24}$$

Proof. The first representation formula is $\Phi\widehat{\Phi}^\star f = f$ for all $f \in W$. Letting $f = W_{on}g$ for some $g \in \mathbb{F}^r$, we have:

$$\begin{aligned}
\Phi\widehat{\Phi}^\star f &= \Phi(W_{on}S_W^{-1}W_{on}^\star\Phi)^\star f \\
&= \Phi\Phi^\star W_{on}(S_W^{-1})^\star W_{on}^\star(W_{on}g) \\
&= SW_{on}S_W^{-1}g \\
&= SW_{on}(W_{on}^\star SW_{on})^{-1}g \qquad (1.25)
\end{aligned}$$

Since $W_{on}W_{on}^\star$ is the identity on W,

$$\begin{aligned}
(1.25) &= W_{on}W_{on}^\star SW_{on}(W_{on}^\star SW_{on})^{-1}g \\
&= W_{on}g \\
&= f
\end{aligned}$$

The second representation formula is $\widehat{\Phi}\Phi^\star f = f$ for all $f \in W$.

$$\begin{aligned}
\widehat{\Phi}\Phi^\star f &= (W_{on}S_W^{-1}W_{on}^\star\Phi)\Phi^\star f \\
&= W_{on}(W_{on}^\star S W_{on})^{-1}W_{on}^\star S W_{on} g \\
&= W_{on} g \\
&= f
\end{aligned}$$

$\square$

The following commutative diagram illustrates the above ideas:

Figure 1.4: Subspace frames diagram

$$\begin{tikzcd}
\boxed{\begin{array}{c}\Phi \text{ subspace frame for } W \subset \mathbb{F}^d \\ (\text{subspace FUNTF for } W \subset \mathbb{F}^d)\end{array}} \arrow[rr, dashed, "W_{on}S_W^{-1}W_{on}^\star"] \arrow[d, "W_{on}^\star"'] & & \boxed{\begin{array}{c}\widehat{\Phi} \text{ subspace frame for } W \subset \mathbb{F}^d \\ (\text{subspace FUNTF for } W \subset \mathbb{F}^d)\end{array}} \\
\boxed{\begin{array}{c}\Phi_W \text{ frame for } \mathbb{F}^r \\ (\text{FUNTF for } \mathbb{F}^r)\end{array}} \arrow[rr, "S_W^{-1}"] & & \boxed{\begin{array}{c}\widehat{\Phi}_W \text{ frame for } \mathbb{F}^r \\ (\text{FUNTF for } \mathbb{F}^r)\end{array}} \arrow[u, "W_{on}"']
\end{tikzcd}$$

We can also extend the frame potential to subspace frames via the following theorem.

Theorem 1.4.6. *For a given s and d, let W be a subspace of $\mathbb{F}^d$ of dimension $r < d$ and consider the resctricted frame potential:*

$$\mathrm{FP}|_W : W \cap \underbrace{(\mathbb{S}^{N-1} \times \cdots \times \mathbb{S}^{N-1})}_{s\ times} \to [0, \infty). \tag{1.26}$$

Then:

1. *Every local minimizer of the restricted frame potential is also a global minimizer.*

2. *If $s \leq r$, the minimum value of the restricted frame potential is s, and the minimizers are precisely the orthonormal sequences in W.*

3. *If $s \geq r$, the minimum value of the restricted frame potential is s^2/r, and the minimizer are precisely the subspace FUNTFs for W.*

Proof. Let W_{on} be a set of r orthonormal vectors such that $\text{span}(W_{on}) = W$ and consider it as an $d \times r$ matrix. If $\Phi = \{\varphi_i : i = 1, \ldots, s\}$ is a finite unit norm set of vectors in W, then the coordinates of Φ in W_{on} are given by the $r \times s$ matrix $\Phi_W = W_{on}^{\star}\Phi$. In [8] it is shown that $\text{FP}(\Phi) = \text{Tr}(S^2)$, where S is the frame operator of Φ. Using the previous two statements we then have:

$$\begin{aligned} \text{FP}|_W(\Phi) &= \text{Tr}(S^2) \\ &= \text{Tr}([(W_{on}\Phi_W)(W_{on}\Phi_W)^{\star}]^2) \\ &= \text{Tr}([\Phi_W\Phi_W^{\star}]^2) \\ &= \text{Tr}(S_W^2) \\ &= \text{FP}(\Phi_W) \end{aligned}$$

Since Φ_W is a finite unit norm set of vectors in $\mathbb{C}^r$, we can apply theorem 1.3.1 to get (1) and (2). Combining theorem 1.3.1 along with theorem 1.4.2 gives (3). $\square$

Chapter 2

Enumeration of Prime Order Harmonic Frames

2.1 Introduction

2.1.1 Harmonic Frames

Harmonic frames are class of FUNTFs that have their origin in the Discrete Fourier Transform (DFT) matrix. The un-normalized $s \times s$ DFT matrix is defined as

$$D_s := (e^{2\pi imn/s})_{m,n=0}^{s-1}. \tag{2.1}$$

Noting that $e^{2\pi imn/s} = e^{2\pi i(m+js)(n+ks)/s}$ for any $j, k \in \mathbb{Z}$, we introduce the additive group of integers mod s,

$$\mathbb{Z}_s = \mathbb{Z}/s\mathbb{Z} := \{0, \ldots, s-1 \bmod s\}. \tag{2.2}$$

Choosing $d \leq s$ distinct columns of D_s, say $n_1, \ldots, n_d \in \mathbb{Z}_s$, we can form the following s vectors in $\mathbb{C}^d$:

$$\varphi_m = \frac{1}{\sqrt{d}}(e^{2\pi imn_j/s})_{j=1}^{d}, \quad m \in \mathbb{Z}_s. \tag{2.3}$$

The set $\Phi = \{\varphi_m : m \in \mathbb{Z}_s\}$ is in fact a FUNTF for $\mathbb{C}^d$, and any frame of this type is called a *DFT-FUNTF*. As we shall see, the DFT-FUNTFs are a subset of the harmonic frames.

Remark 2.1.1. Since we will be dealing exclusively with finite frames in this chapter, we shall interchangeably consider the frame Φ as a set or a matrix (whose columns are the vectors φ_i), with the appropriate usage being determined by the context. See section 1.3 (Finite frame theory) for details on considering a frame as a matrix.

Let G denote a group. Define $\mathbb{C}^\times$ as the group of units of $\mathbb{C}$, that is the set $\mathbb{C}\backslash\{0\}$ endowed with multiplication. A *character* of a group G is a group homomorphism $\xi : G \to \mathbb{C}^\times$ that satisfies

$$\xi(gg') = \xi(g)\xi(g'), \quad \forall\, g, g' \in G. \tag{2.4}$$

Suppose G is a finite group of order s, i.e.

$$G = \{g_i : i = 1, \ldots, s\}. \tag{2.5}$$

Then for each $g_i \in G$, $\xi(g_i)$ is a s-th root of unity. If G is also abelian, then it has exactly s characters, $\{\xi_i : i = 1, \ldots, s\}$. The set of vectors $\{(\xi_i(g_j))_{j=1}^s : i = 1, \ldots, s\}$ form an orthogonal basis for $\mathbb{C}^s$. The matrix with these vectors as rows,

$$(\xi_i(g_j))_{i,j=1}^s, \tag{2.6}$$

is the *character table* of G. In particular, when $G \cong \mathbb{Z}_s$, the character table of G is D_s.

Let $\mathcal{U}(\mathbb{C}^d)$ denote the group of unitary transformations on $\mathbb{C}^d$, i.e.

$$\mathcal{U}(\mathbb{C}^d) := \{U \in \mathcal{M}_{d\times d}(\mathbb{C}) : U^\star U = UU^\star = I\}. \tag{2.7}$$

Furthermore, let $\mathcal{I} \subseteq \{1, \ldots, s\}$ with $|\mathcal{I}| = d$, and suppose G is a finite abelian group of order s. Then, for any $U \in \mathcal{U}(\mathbb{C}^d)$, the set,

$$\Phi = \{U(\xi_i(g_j))_{i\in\mathcal{I}} : j = 1, \ldots, s\} \tag{2.8}$$

is a frame for $\mathbb{C}^d$ and is called a *harmonic frame*. Note that when $G \cong \mathbb{Z}_s$ and $U = I$, one obtains a DFT-FUNTF.

Important in the study of harmonic frames is the notion of the symmetry group. The *symmetry group* of a FUNTF Φ for $\mathbb{C}^d$ is the group:

$$\mathrm{Sym}(\Phi) := \{U \in \mathcal{U}(\mathbb{C}^d) : \{U\varphi_i : i = 1, \ldots, s\} = \{\varphi_i : i = 1, \ldots, s\}\}. \tag{2.9}$$

We can recast the definition of symmetry group in terms matrices. Let S_k denote the group of permutations on k elements. We say $P \in \mathcal{M}_{d\times d}(\mathbb{C})$ is a *permutation matrix* if there exists a permutation $\sigma \in S_d$ such that

$$P_{i,j} = \begin{cases} 1, & \text{if } j = \sigma(i) \\ 0, & \text{otherwise.} \end{cases} \tag{2.10}$$

Let $\mathcal{P}_s$ denote the set of all $s \times s$ permuation matrices. Then, in terms of matrices, the symmetry group of Φ is

$$\mathrm{Sym}(\Phi) = \{U \in \mathcal{U}(\mathbb{C}^d) : \exists\, P \in \mathcal{P}_s \text{ such that } U\Phi = \Phi P\}. \tag{2.11}$$

While there has been much work on harmonic frames and subjects related to them (see, for example, [20, 27, 33, 34, 36]), we will need only the following result from [33].

Theorem 2.1.2. *A FUNTF Φ of s vectors for $\mathbb{C}^d$ is harmonic if and only if it is generated by a finite abelian group $G \subset \mathrm{Sym}(\Phi)$ of order s, i.e., $\Phi = G\varphi$, for all $\varphi \in \Phi$.*

2.1.2 The Enumeration Problem

The purpose of this chapter is to count all equivalence classes of prime order harmonic frames. The definition of what it means for two harmonic frames to be equivalent will be given in section 2.2. We start with simpler problem concerning the enumeration of DFT-FUNTFs.

Recall the definition of a DFT-FUNTF given by equation (2.3). A basic way of counting the number of DFT-FUNTFs is inspired by the following observation. For any vector $f \in \mathbb{C}^d$, the frame Φ gives the following representation of f:

$$f \mapsto (\langle f, \varphi_m \rangle)_{m=0}^{s-1} \in \mathbb{C}^s. \tag{2.12}$$

Therefore, even a re-indexing of the frame would change the representation it gives for a fixed f. Thus, we could count the number of ordered DFT-FUNTFs. To accomplish this task, we observe that there are s columns in D_s and we select d of them. Since each ordered combination of column choices $n_1, \ldots, n_d$ gives a distinct frame, there are $s(s-1)\cdots(s-d+1)$ ordered DFT-FUNTFs.

There are of course other ways by which we may distinguish frames, and we shall consider two others here. The first is a natural counterpart to the ordered counting scheme, namely, counting the number of DFT-FUNTFs considered as unordered sets of vectors. The techniques developed for this method will then be expanded to our main goal, which is to count all inequivalent harmonic frames of prime order, where two harmonic frames shall be considered equivalent if one is the unitary transformation of the other. As we shall see, this amounts to counting the number of inequivalent DFT-FUNTFs.

There has been some interest in harmonic frames in the literature, see [20, 33]. In particular, [34] presents a computer program for generating all equivalence classes of harmonic frames for a given s and d, where there is a limit on the size of either due to computational considerations. From this program, the authors conjecture that there are $\mathcal{O}(s^{d-1})$ inequivalent harmonic frames. The content of this chapter is to not only validate this conjecture for the case when s is a prime number, but in fact give an exact formula for the number of harmonic frames in this case. Furthermore, we examine the structure of prime order harmonic frames via their symmetry group.

An outline of the remainder of chapter 2 is as follows: the rest of section 2.1 reviews some algebraic theory and examines the problem of counting unordered DFT-FUNTFs. Section 2.2 presents the main result of this paper. In section 2.3 we define an equivalence relation that is equivalent to (2.23) and then use this to develop a correspondence between inequivalent harmonic frames and the orbits of a particular set. Section 2.4 counts the number of orbits of this particular set, thus giving a formula for the number of inequivalent harmonic frames. The structure of the symmetry group is handled in section 2.5, and section 2.6 contains a few concluding remarks.

2.1.3 Algebra Review[1]

Recall that we denote the additive group of integers mod s by $\mathbb{Z}_s$, and set

$$\mathbb{Z}_s^d := \underbrace{\mathbb{Z}_s \times \cdots \times \mathbb{Z}_s}_{d \text{ times}}. \tag{2.13}$$

[1] All material in this section can be found in [19]

Furthermore, let $\mathbb{Z}_s^\times$ denote the group of units of $\mathbb{Z}_s$, which, when s is prime, is simply the set $\{1, \ldots, s\}$ endowed with multiplication mod s. Finally, for $k \in \mathbb{N}$, let S_k denote the group of permutations of k elements. We will also need the following definitions and proposition:

Definition 2.1.3. A *group action* of a group G on a set X is a map π,

$$\begin{aligned} \pi : G \times X &\to X \\ (g, x) &\mapsto g \cdot x, \end{aligned}$$

satisfying the following properties:

1) $g_1 \cdot (g_2 \cdot x) = (g_1 g_2) \cdot x \ \forall \ g_1, g_2 \in G, \ x \in X,$

2) $1 \cdot x = x \ \forall \ x \in X.$

Definition 2.1.4. Let X be some set and let G be a group. Furthermore, let $\pi : G \times X \to X$ be a group action. For each $x \in X$ the *stabilizer* of x in G is the subgroup of G that fixes the element x:

$$G_x := \{g \in G : g \cdot x = x\}. \tag{2.14}$$

Proposition 2.1.5. *Let G be a group acting on the nonempty set X. The relation on X defined by:*

$$x_1 \sim x_2 \iff x_1 = g \cdot x_2 \text{ for some } g \in G$$

is an equivalence relation. For each $x \in X$, the number of elements in the equivalence class containing x is $|G : G_x|$, the index of the stabilizer of x.

Note, when G is a finite group,

$$|G : G_x| = \frac{|G|}{|G_x|}. \tag{2.15}$$

Definition 2.1.6. Let G be a group acting on the nonempty set X. The equivalence class $\mathcal{O}_x := \{g \cdot x : g \in G\}$ is called the *orbit* of G containing x.

As such, the orbits of a group action disjointly partition the set X. We are now ready to count the number of prime order DFT-FUNTFs, considered as unordered sets. The basic structure of the argument in subsection 2.1.4 will be used when we count all harmonic frames of prime order, albeit with added complexity.

2.1.4 The Number of Unordered DFT-FUNTFs

It is often the case that we would like to consider a frame as a set, where the order of elements does not matter. Given two ordered DFT-FUNTFs $\Phi = (\varphi_0, \ldots, \varphi_{s-1})$ and $\Psi = (\psi_0, \ldots, \psi_{s-1})$, we define the following equivalence relation:

$$\Phi \sim_1 \Psi \iff \exists\, \sigma \in S_s \;\; s.t. \;\; \varphi_m = \psi_{\sigma(m)}, \quad \forall \; m = 0, \ldots, s-1. \tag{2.16}$$

(2.16) merely formalizes our consideration of frames as sets. An equivalence class of (2.16) will be denoted in the usual way, that is $\Phi = \{\varphi_0, \ldots, \varphi_{s-1}\}$. In this subsection, we count the number of DFT-FUNTFs of prime order under (2.16). First, however, we must change our perspective on the problem.

Remark 2.1.7. For the rest of chapter 2 we will only consider unordered DFT-FUNTFs, and as such from now on Φ will denote $\{\varphi_0, \ldots, \varphi_{s-1}\}$.

2.1.4.1 DFT-FUNTFs and Orbits

First notice that every DFT-FUNTF contains the vector $\varphi_0 = \frac{1}{\sqrt{d}}(1, \ldots, 1) \in \mathbb{C}^d$, and so when comparing two such frames we need not consider this vector. Thus we will only compare sets of the form

$$\Phi' = \Phi - \{\varphi_0\}. \tag{2.17}$$

Define the set $\tilde{\mathbb{Z}}_s^d$ as

$$\tilde{\mathbb{Z}}_s^d := \{n = (n_1, \ldots, n_d) \in \mathbb{Z}_s^d : n_i \neq n_j, \ \ \forall \ i \neq j\}. \tag{2.18}$$

There is a one-to-one correspondence between the vectors φ_m, $m \neq 0$, and the elements of $\tilde{\mathbb{Z}}_s^d$. Considering $\mathbb{Z}_s^\times$ as a group and $\tilde{\mathbb{Z}}_s^d$ as a set, we define the group action π_1 as:

$$\begin{aligned} \pi_1 : \mathbb{Z}_s^\times \times \tilde{\mathbb{Z}}_s^d &\to \tilde{\mathbb{Z}}_s^d \\ (m, n) &\mapsto m \cdot n := (mn_1, \ldots, mn_d). \end{aligned} \tag{2.19}$$

The orbits of π_1 are then the sets

$$\mathcal{O}_n = \{m \cdot n = (mn_1, \ldots, mn_d) : m \in \mathbb{Z}_s^\times\}, \quad n \in \tilde{\mathbb{Z}}_s^d. \tag{2.20}$$

Remark 2.1.8. For clarity of exposition we shall sometimes use Φ_n to denote the DFT-FUNTF Φ and $\varphi_{m,n}$ its corresponding elements, where the subscript n emphasizes the generators $n = (n_1, \ldots, n_d)$.

The following proposition relates the equivalence classes of (2.16) and the orbits of π_1.

Proposition 2.1.9. *There is a one-to-one correspondence between the equivalence classes of (2.16) and the orbits of* π_1*, i.e. the sets* Φ_n *and* $\mathcal{O}_n$ *can be identified. We denote this identification as:*

$$\Phi_n = \{\varphi_0, \ldots, \varphi_{s-1}\} \longleftrightarrow \mathcal{O}_n. \tag{2.21}$$

Proof. As noted above, we have:

$$\Phi \longleftrightarrow \Phi' = \Phi - \{\varphi_0\}.$$

Define a function F that maps orbits of $\tilde{\mathbb{Z}}_s^d$ to sets of the form Φ' as follows:

$$F(\mathcal{O}_n) = \{\varphi_{m,n}\}_{m=1}^{s}. \tag{2.22}$$

We must show that F is both one-to-one and onto, however it is clear that F is surjective. Considering then the former, suppose $F(\mathcal{O}_n) = F(\mathcal{O}_{n'})$. This would imply that $\{\varphi_{m,n}\}_{m=1}^{s} = \{\varphi_{m',n'}\}_{m'=1}^{s}$. But then for some m and some m', we would have $(mn_1, \ldots, mn_d) = (m'n'_1, \ldots, m'n'_d)$, i.e. $\mathcal{O}_n \cap \mathcal{O}_{n'} \neq \emptyset$, and so in fact $\mathcal{O}_n = \mathcal{O}_{n'}$. □

Remark 2.1.10. Given the content of proposition 2.1.9, we now replace the problem of counting the equivalence classes of (2.16) with the problem of counting the orbits of π_1.

2.1.4.2 The Number of Orbits of π_1

By proposition 2.1.5 we see that the orbits of a group action partition the set into disjoint equivalence classes. In particular, the orbits $\mathcal{O}_n$ partition the set $\tilde{\mathbb{Z}}_s^d$.

Furthermore, the size of each $\mathcal{O}_n$ is given by $|\mathcal{O}_n| = |\mathbb{Z}_s^\times : (\mathbb{Z}_s^\times)_n|$. Using these facts, we prove the following proposition.

Proposition 2.1.11. *Let s be a prime number and $d \leq s$. Then the number of orbits of π_1 is:*

1) 2, *if $d = 1$ or $d = s = 2$.*

2) $s(s-2)\cdots(s-d+1)$, *if $d \geq 2$, $s > 2$.*

Proof. We first consider the case $d = 1$. For $n = 0$ we have $(\mathbb{Z}_s^\times)_0 = \mathbb{Z}_s^\times$, and so $|\mathcal{O}_0| = (s-1)/(s-1) = 1$. For $n \neq 0$ we see $(\mathbb{Z}_s^\times)_n = \{1\}$, and thus $|\mathcal{O}_n| = s - 1$. Since $|\tilde{\mathbb{Z}}_s^1| = s$, there are only two orbits.

Now take $2 \leq d \leq s$. For each $n \in \tilde{\mathbb{Z}}_s^d$ we have $(\mathbb{Z}_s^\times)_n = \{1\}$, and thus $|\mathcal{O}_n| = s - 1$. Therefore the number of orbits is given by γ, where

$$\begin{aligned} |\tilde{\mathbb{Z}}_s^d| &= \gamma|\mathcal{O}_n|, \\ s(s-1)\cdots(s-d+1) &= \gamma(s-1). \end{aligned}$$

For $s = 2$ and $d = 2$, we see $\gamma = 2$. For $s > 2$ we have $\gamma = s(s-2)\cdots(s-d+1)$. $\square$

As an addendum to theorem 2.1.11, we note that one of the orbits in the $d = 1$ case corresponds to a degenerate DFT-FUNTF. Namely, the orbit $\mathcal{O}_0$ corresponds to the DFT-FUNTF consisting of the single element $\{1\}$.

2.2 The Number of Harmonic Frames of Prime Order

Using a similar correspondence between harmonic frames and orbits, we count all harmonic frames of prime order up to unitary transformations.

Two harmonic frames $\Phi = \{\varphi_0, \ldots, \varphi_{s-1}\} \subset \mathbb{C}^d$ and $\Psi = \{\psi_0, \ldots, \psi_{s-1}\} \subset \mathbb{C}^d$ are said to be *equivalent* if the following equivalence relation holds:

$$\Phi \sim_2 \Psi \iff \exists\ U \in \mathcal{U}(\mathbb{C}^d) \text{ and } P \in \mathcal{P}_s \text{ s.t. } \Phi = U\Psi P. \tag{2.23}$$

Note that we have used matrix notation for the left hand side of (2.23). In terms of sets, the condition merely states that

$$\{\varphi_i : i = 1, \ldots s\} = \{U\psi_i : i = 1, \ldots, s\}. \tag{2.24}$$

(2.23) is a standard form of equivalence in much of the literature when dealing with frames. Recently, [34] conjectured that the number of *inequivalent* harmonic frames is $O(s^{d-1})$. We prove this conjecture for s a prime number as a corollary to theorem 2.2.1, which gives an exact formula for the number of harmonic frames. The proof of theorem 2.2.1 is handled in section 2.4, with much preliminary work accomplished in section 2.3.

For a fixed s and d, we backwards recursively define the set

$$\{\alpha_c \in \mathbb{N} \cup \{0\} : c \in \mathbb{N}, c \mid s-1, \text{ and } c \mid d \text{ or } c \mid d-1\}. \tag{2.25}$$

If $c \mid s-1$, $c \mid d$, and $c > 1$, then

$$\alpha_c := \frac{(s-1-c)(s-1-2c)\cdots(s-1-(\frac{d}{c}-1)c)}{c^{\frac{d}{c}-1}(d/c)!} - \frac{c}{s-1}\sum_{\substack{c<b<s\\ c|b,\, b|d}} \binom{s-1}{b}\alpha_b, \tag{2.2.26 d}$$

where we have used the notation (2.2.26 d) to emphasize its dependence on the condition $c \mid d$. If $c \mid s-1$, $c \mid d-1$, and $c > 1$, then

$$\alpha_c := \frac{(s-1-c)(s-1-2c)\cdots(s-1-(\frac{d-1}{c}-1)c)}{c^{\frac{d-1}{c}-1}((d-1)/c)!} - \frac{c}{s-1}\sum_{\substack{c<b<s\\ c|b,\, b|d-1}} \binom{s-1}{b}\alpha_b. \tag{2.2.26 $d-1$}$$

Finally, α_1 is defined as:

$$\alpha_1 := \frac{1}{s-1}\binom{s}{d} - \sum_{\substack{c|d \\ c>1}} \frac{\alpha_c}{c} - \sum_{\substack{c|d-1 \\ c>1}} \frac{\alpha_c}{c}. \tag{2.2.27}$$

Theorem 2.2.1. *Let s be a prime number and let $1 < d < s$. Define the set*

$$\{\alpha_c \in \mathbb{N} \cup \{0\} : c \in \mathbb{N}, c \mid s-1, \textit{ and } c \mid d \textit{ or } c \mid d-1\},$$

as in equations (2.2.26 d), (2.2.26 d − 1), and (2.2.27). The total number of harmonic frames for $\mathbb{C}^d$ with s elements is then given by:

$$\alpha_1 + \sum_{\substack{c|d \\ c>1}} \alpha_c + \sum_{\substack{c|d-1 \\ c>1}} \alpha_c. \tag{2.2.28}$$

More concisely, we have the following corollary:

Corollary 2.2.2. *Let s be any prime number and fix d such that $1 < d < s$. Then the number of inequivalent harmonic frames for $\mathbb{C}^d$ with s elements is $O(s^{d-1})$.*

Proof. Using equations (2.2.26 d) and (2.2.26 $d-1$), we see that $\alpha_c = O(s^{d'})$, where $c > 1$ and $d' \leq \frac{d}{c} - 1 < d-1$. Therefore, by (2.2.27), we see that $\alpha_1 = O(s^{d-1})$, and the corollary follows. □

In the above theorems, the case $d = 1$ is omitted, however, it is not hard to see that there are two inequivalent harmonic frames in this case; in fact, there is only one inequivalent harmonic frame for $d = 1$ with s distinct vectors.

2.3 Harmonic Frames and Orbits

In this section we develop a one-to-one correspondence between inequivalent harmonic frames and the orbits of a particular set, not unlike the ideas presented in subsection 2.1.4. First, however, we come up with an equivalent condition to (2.23). We will assume s is prime for the remainder of chapter 2.

2.3.1 A New Equivalence Relation

When s is prime, every harmonic frame is of the form $U\Phi$, where $U \in \mathcal{U}(\mathbb{C}^d)$ and Φ is a DFT-FUNTF (see section 2.1.1). Therefore, finding the number of inequivalent harmonic frames amounts to finding the number of inequivalent DFT-FUNTFs. Toward that end, we simplify (2.23) to the following:

Theorem 2.3.1. *If s is prime and $\Phi = \{\varphi_0, \ldots, \varphi_{s-1}\}$ and $\Psi = \{\psi_0, \ldots, \psi_{s-1}\}$ are DFT-FUNTFs, then*

$$\Phi \sim_2 \Psi \quad \Longleftrightarrow \quad \begin{array}{c} \exists\ \sigma_1 \in S_s,\ \ \sigma_2 \in S_d \textit{ such that} \\ \varphi_m(k) = \psi_{\sigma_1(m)}(\sigma_2(k)) \\ \forall\ m = 0, \ldots, s-1,\ \ k = 1, \ldots, d, \end{array} \tag{2.3.29}$$

where $\varphi_m(k)$ denotes the k^{th} element of the vector φ_m.

Proof. It is clear that if the right hand side of (2.3.29) holds, then the left hand side must hold as well. Assume then that $\Phi \sim_2 \Psi$, i.e., there exists a $U \in \mathcal{U}(\mathbb{C}^d)$ and $P_\sigma \in \mathcal{P}_s$ such that

$$\Phi = U\Psi P_\sigma. \tag{2.3.30}$$

Let $\sigma \in S_s$ be the permutation associated with P_σ, and note that (2.3.30) implies that

$$\varphi_m = U\psi_{\sigma(m)}, \quad \forall\ m = 0, \ldots, s-1. \tag{2.3.31}$$

Without loss of generality, we may assume that $\sigma(0) = 0$. Indeed, by theorem 2.1.2 there exists a $U_0 \in \mathrm{Sym}(\Psi)$ such that $U_0\psi_0 = \psi_{\sigma(0)}$. By definition, U_0 is a $d \times d$ matrix that permutes the columns of Ψ by acting on the left. Therefore, there exists an $s \times s$ permutation matrix P_{U_0} that permutes the columns of Ψ in the exact same manner, yet acts on the right. In particular, $U_0\Psi = \Psi P_{U_0}$, and thus

$$\Phi = UU_0\Psi P_{U_0}^{-1} P_\sigma. \tag{2.3.32}$$

Set $V := UU_0$ and $P := P_{U_0}^{-1} P_\sigma$. It is clear that V is a unitary transformation and that P is its associated permutation matrix. Furthermore, $\varphi_0 = V\psi_0$, and so we can assume from the start that $\varphi_0 = U\psi_0$, i.e., that $\sigma(0) = 0$.

Now let $n_1, \ldots, n_d$ denote the column choices of Φ, and consider the following:

$$\langle \varphi_m, \varphi_0 \rangle = \sum_{k=1}^{d} e^{2\pi i m n_k/s}. \tag{2.3.33}$$

Letting $l_1, \ldots, l_d$ denote the column choices of Ψ, we also have:

$$\langle \varphi_m, \varphi_0 \rangle = \langle U\psi_{\sigma(m)}, U\psi_0 \rangle = \langle \psi_{\sigma(m)}, \psi_0 \rangle = \sum_{k=1}^{d} e^{2\pi i \sigma(m) l_k/s}. \tag{2.3.34}$$

Define $p_\varphi, p_\psi \in \mathbb{Z}[z]/\langle z^s \rangle$ as follows:

$$p_\varphi(z) := \sum_{k=1}^{d} z^{mn_k} \quad \text{and} \quad p_\psi(z) := \sum_{k=1}^{d} z^{\sigma(m) l_k}. \tag{2.3.35}$$

By equations (2.3.33) and (2.3.34), we see that $p_\varphi(z) = p_\psi(z)$ when $z = e^{2\pi i/s}$. In other words, $z = e^{2\pi i/s}$ is a root of the polynomial $p(z) := p_\varphi(z) - p_\psi(z)$. However, since $p \in \mathbb{Z}[z]/\langle z^s \rangle$, and the minimum polynomial of $z = e^{2\pi i/s}$ is $q(z) := \sum_{k=0}^{s-1} z^k$, p must either be an integer multiple of q or the zero polynomial. It is clear, though, that only the latter option is feasible, thus giving

$$p_\varphi(z) = p_\psi(z). \tag{2.3.36}$$

Combining equations (2.3.35) and (2.3.36), we see there exists a $\sigma_2 \in S_d$ such that

$$mn_k = \sigma(m) l_{\sigma_2(k)}, \quad \forall \ k = 1, \ldots, d. \tag{2.3.37}$$

Note that σ_2 is dependent on the choice of m. Taking $m = 1$ in (2.3.37), one has $n_k = \sigma(1) l_{\sigma_2(k)}$. Letting $\sigma_1(m) := \sigma(1)m$, we have:

$$\varphi_m = (e^{2\pi i m n_k / s})_{k=1}^d = (e^{2\pi i \sigma_1(m) l_{\sigma_2(k)}})_{k=1}^d = \psi_{\sigma_1(m)}(\sigma_2(k)). \tag{2.3.38}$$

$\square$

2.3.2 Inequivalent DFT-FUNTFs and Orbits

Similar to section 2.1.4.1, we now develop a one-to-one correspondence between inequivalent DFT-FUNTFs and the orbits of a particular set. As a matter of notation, we shall denote equivalence classes of (2.23) by $[\Phi]$, where $\Phi = \{\varphi_0, \ldots, \varphi_{s-1}\}$ is a DFT-FUNTF representative. By theorem 2.3.1, the equivalence classes of (2.23) are identical to the equivalence classes of the right hand side of (2.3.29). We now turn our attention to the set with which we will identify the equivalence classes $[\Phi]$.

Consider the following equivalence relation on the set $\tilde{\mathbb{Z}}_s^d$,

$$(n_1, \ldots, n_d) \sim (n_1', \ldots, n_d') \iff \exists \ \sigma \in S_d \ \ s.t. \ (n_1, \ldots, n_d) = (n_{\sigma(1)}', \ldots, n_{\sigma(d)}'). \tag{2.3.39}$$

Denote an equivalence class of (2.3.39) by the representative $[n] = [n_1, \ldots, n_d]$, and define $\mathbb{A}_s^d$ as the set of all equivalence classes, i.e.

$$\mathbb{A}_s^d := \tilde{\mathbb{Z}}_s^d / \sim . \tag{2.3.40}$$

It is easy to see $|\mathbb{A}_s^d| = \binom{s}{d}$. Considering $\mathbb{Z}_s^\times$ as a group and $\mathbb{A}_s^d$ as a set, we define the group action π_2,

$$\begin{aligned} \pi_2 : \mathbb{Z}_s^\times \times \mathbb{A}_s^d &\rightarrow \mathbb{A}_s^d \\ (m, [n]) &\mapsto m \cdot [n] := [mn_1, \ldots, mn_d]. \end{aligned} \tag{2.3.41}$$

The orbits of π_2 are the sets $\mathcal{O}_{[n]} = \{m \cdot [n] = [mn_1, \ldots, mn_d] : m \in \mathbb{Z}_s^\times\}$. The following proposition relates the equivalence classes of (2.23) and the orbits of π_2.

Proposition 2.3.2. *There is a one-to-one correspondence between the equivalences classes of* (2.23) *and the orbits of* π_2*, i.e.*

$$[\Phi_n] \longleftrightarrow \mathcal{O}_{[n]}. \tag{2.3.42}$$

Proof. Define the function F as follows:

$$F([\Phi_n]) = \mathcal{O}_{[n]} = \{[mn_1, \ldots, mn_d] : m \in \mathbb{Z}_s^\times\}. \tag{2.3.43}$$

We must show that F is well defined, one-to-one, and onto. Surjectivity is clear, so we focus on the first two. To show F is well defined, suppose that $[\Phi_n] = [\Psi_{n'}]$. We want to show $F([\Phi_n]) = F([\Psi_{n'}])$, i.e. $\mathcal{O}_{[n]} = \mathcal{O}_{[n']}$. We have:

$$\begin{aligned} [\Phi_n] = [\Psi_{n'}] \quad &\Longleftrightarrow \quad \varphi_m(k) = \psi_{\sigma_1(m)}(\sigma_2(k)) \ \forall \ k = 1, \ldots, d, \ \forall \ m = 0, \ldots, s-1 \\ &\Longleftrightarrow \quad \{\varphi_0(k)_{k=1}^d, \ldots, \varphi_{s-1}(k)_{k=1}^d\} = \{\psi_0(\sigma_2(k))_{k=1}^d, \ldots, \psi_{s-1}(\sigma_2(k))_{k=1}^d\} \\ &\Longleftrightarrow \quad \{\varphi_1(k)_{k=1}^d, \ldots, \varphi_{s-1}(k)_{k=1}^d\} = \{\psi_1(\sigma_2(k))_{k=1}^d, \ldots, \psi_{s-1}(\sigma_2(k))_{k=1}^d\} \\ &\Longleftrightarrow \quad \{(mn_1, \ldots, mn_d) : m \in \mathbb{Z}_s^\times\} = \{(mn'_{\sigma_2(1)}, \ldots, mn'_{\sigma_2(d)}) : m \in \mathbb{Z}_s^\times\} \\ &\Longleftrightarrow \quad \{[mn_1, \ldots, mn_d] : m \in \mathbb{Z}_s^\times\} = \{[mn'_1, \ldots, mn'_d] : m \in \mathbb{Z}_s^\times\} \\ &\Longleftrightarrow \quad \mathcal{O}_{[n]} = \mathcal{O}_{[n']}, \end{aligned}$$

where the first equivalence is due to theorem 2.3.1, and the third equivalence is because $\varphi_0 = \psi_0 = \frac{1}{\sqrt{d}}(1, \ldots, 1)$.

To prove injectivity, we assume $\mathcal{O}_{[n]} = \mathcal{O}_{[n']}$. According to this assumption, there must exist an $m_0' \in \mathbb{Z}_s^\times$ such that $[n_1, \ldots, n_d] = [m_0'n_1', \ldots, m_0'n_d']$. Therefore we have:

$$\begin{aligned}
\mathcal{O}_{[n]} = \mathcal{O}_{[n']} \quad &\iff \quad [n_1, \ldots, n_d] = [m_0'n_1', \ldots, m_0'n_d'] \\
&\iff \quad (n_1, \ldots, n_d) = (m_0'n'_{\sigma_2(1)}, \ldots, m_0'n'_{\sigma_2(d)}) \\
&\iff \quad (mn_1, \ldots, mn_d) = (mm_0'n'_{\sigma_2(1)}, \ldots, mm_0'n'_{\sigma_2(d)}), \ \ \forall \ m \in \mathbb{Z}_s^\times \\
&\iff \quad \{(mn_1, \ldots, mn_d) : m \in \mathbb{Z}_s^\times\} = \{(mn'_{\sigma_2(1)}, \ldots, mn'_{\sigma_2(d)}) : m \in \mathbb{Z}_s^\times\} \\
&\iff \quad \{\varphi_1(k)_{k=1}^d, \ldots, \varphi_{s-1}(k)_{k=1}^d\} = \{\psi_1(\sigma_2(k))_{k=1}^d, \ldots, \psi_{s-1}(\sigma_2(k))_{k=1}^d\} \\
&\iff \quad \{\varphi_0(k)_{k=1}^d, \ldots, \varphi_{s-1}(k)_{k=1}^d\} = \{\psi_0(\sigma_2(k))_{k=1}^d, \ldots, \psi_{s-1}(\sigma_2(k))_{k=1}^d\} \\
&\iff \quad \varphi_m(k) = \psi_{\sigma_1(m)}(\sigma_2(k)), \ \ \forall \ k = 1, \ldots, d, \ \ m = 0, \ldots, s-1 \\
&\iff \quad [\Phi_n] = [\Psi_{n'}],
\end{aligned}$$

where the fourth equivalence uses the fact that $\{mm_0' : m \in \mathbb{Z}_s^\times\} = \{m : m \in \mathbb{Z}_s^\times\}$. $\square$

To conclude this section, we note that when $d = s$, we see $|\mathbb{A}_s^d| = 1$, and so there can be only one orbit. Thus there is only one harmonic frame in this case.

2.4 The Number of Orbits of $\mathbb{A}_s^d$

We begin by counting the number of orbits of $\mathbb{A}_s^d$ under the group action π_2 for the cases $d = 2$ and $d = 3$. We then generalize these results for all $1 < d < s$.

2.4.1 Some Examples: d = 2 and d = 3

Proposition 2.4.1. *Let s be an odd prime number and let $d = 2$. Then there are $(s+1)/2$ orbits of $\mathbb{A}_s^2$. Therefore, there are $(s+1)/2$ inequivalent harmonic frames for $\mathbb{C}^2$.*

Proof. Let $[n] \in \mathbb{A}_s^2$. If $(\mathbb{Z}_s^\times)_{[n]} = \{1\}$, then $|\mathcal{O}_{[n]}| = s-1$. Therefore, if we can find all $[n] \in \mathbb{A}_s^2$ with non-trivial stabilizer and their corresponding orbits, we will be able to solve for the total number of orbits. Assume that $m \cdot [n_1, n_2] = [mn_1, mn_2] = [n_1, n_2]$ for some $m \neq 1$. This implies that

$$\begin{aligned} mn_1 &\equiv n_2 \bmod s, \\ mn_2 &\equiv n_1 \bmod s. \end{aligned}$$

Combining the above equations yields

$$\begin{aligned} m^2 n_1 &\equiv n_1 \bmod s \\ \Rightarrow \qquad m &\equiv \pm 1 \bmod s. \end{aligned}$$

Thus we see that we can take $m \equiv -1 \bmod s$, which implies $n_2 \equiv -n_1 \bmod s$. Therefore all $[n] \in \mathbb{A}_s^2$ of the form $[n] = [n_1, -n_1]$, $n_1 \neq 0$, have stabilizer $\{1, -1\}$. Furthermore, since

$$\mathcal{O}_{[1,-1]} = \{m \cdot [1, -1] = [m, -m] : m \in \mathbb{Z}_s^\times\}, \tag{2.4.44}$$

we see that all such $[n]$ lie in the orbit $\mathcal{O}_{[1,-1]}$. Finally, these are the only elements of $\mathbb{A}_s^2$ with nontrivial stabilizer, and thus the number of orbits of $\mathbb{A}_s^2$ is $\gamma_1 + 1$, where

γ_1 is the number of orbits of size $s-1$. Therefore,

$$\begin{aligned}
|\mathbb{A}_s^2| &= \gamma_1(s-1) + |\mathcal{O}_{(1,-1)}|, \\
\binom{s}{2} &= \gamma_1(s-1) + (s-1)/2, \\
s(s-1)/2 &= \gamma_1(s-1) + (s-1)/2.
\end{aligned}$$

Solving for γ_1 we get $\gamma_1 = (s-1)/2$ and so $\mathbb{A}_s^2$ has $\gamma_1 + 1 = (s-1)/2 + 1 = (s+1)/2$ orbits. □

Proposition 2.4.2. *Let s be a prime number, $s > 3$, and let $d = 3$:*

1. *If $s \equiv 1 \mod 3$, then there are $(s^2 - 2s + 7)/6$ orbits of $\mathbb{A}_s^3$.*

2. *If $s \equiv 2 \mod 3$, then there are $(s^2 - 2s + 3)/6$ orbits of $\mathbb{A}_s^3$.*

Therefore, if $s \equiv 1 \mod 3$, there are $(s^2 - 2s + 7)/6$ inequivalent harmonic frames for $\mathbb{C}^3$, and if $s \equiv 2 \mod 3$, there are $(s^2 - 2s + 3)/6$ inequivalent harmonic frames for $\mathbb{C}^3$.

Proof. As in the proof of proposition 2.4.1, we are looking for all $[n] \in \mathbb{A}_s^3$ with non-trivial stabilizer and their corresponding orbits. So again we suppose

$$m \cdot [n_1, n_2, n_3] = [mn_1, mn_2, mn_3] = [n_1, n_2, n_3], \tag{2.4.45}$$

for some $m \neq 1$. We now consider two cases:

I: Suppose $n_1 = 0$. Then we want $m \cdot [0, n_2, n_3] = [0, mn_2, mn_3] = [0, n_2, n_3]$. But this is just the same situation as the $d = 2$ case, and so the elements of $\mathbb{A}_s^3$ of this form with non-trivial stabilizer all lie in the following orbit:

$$\begin{aligned}
\mathcal{O}_{[0,1,-1]} &= \{m \cdot [0,1,-1] = [0, m, -m] : m \in \mathbb{Z}_s^\times\}, \\
|\mathcal{O}_{[0,1,-1]}| &= (s-1)/2.
\end{aligned} \tag{2.4.46}$$

II: Suppose $n_k \neq 0$ for all $k = 1, 2, 3$. According to (2.4.45), we have three options for the value of mn_1:

$$mn_1 \equiv \begin{cases} n_1 \bmod s, \\ n_2 \bmod s, \\ n_3 \bmod s. \end{cases} \tag{2.4.47}$$

If $mn_1 \equiv n_1 \bmod s$, then $m = 1$ is the only solution, which is trivial and so we disregard this case. Since the order of elements does not matter in $\mathbb{A}_s^3$, there is no difference between $mn_1 \equiv n_2 \bmod s$ and $mn_1 \equiv n_3 \bmod s$, and so we choose the former. Moving on to the value of mn_2, we once again have the same three options. However, $mn_2 \equiv n_1 \bmod s$, combined with $mn_1 \equiv n_2 \bmod s$ would imply that $mn_3 \equiv n_3 \bmod s$, thus resulting in $m = 1$. $mn_2 \equiv n_2 \bmod s$ not only would imply $m = 1$, but since $mn_1 \equiv n_2 \bmod s$, would also lead to a contradiction. Therefore $mn_2 \equiv n_3 \bmod s$ must hold, which in turn forces $mn_3 \equiv n_1 \bmod s$. Summarizing, we have

$$\begin{aligned} mn_1 &\equiv n_2 \bmod s, \\ mn_2 &\equiv n_3 \bmod s, \\ mn_3 &\equiv n_1 \bmod s. \end{aligned} \tag{2.4.48}$$

Proceeding in a similar fashion to the proof of proposition 2.4.1, we see that (2.4.48) implies

$$m^3 n_1 \equiv n_1 \bmod s. \tag{2.4.49}$$

We now find all $m \in \mathbb{Z}_s^\times$ that satisfy (2.4.49). Naturally $m = 1$ works; for the remaining solutions, let g be any primitive root mod s, i.e. $\langle g \rangle = \mathbb{Z}_s^\times$. Then all nontrivial solutions to (2.4.49) are of the form

$$m \equiv g^{(s-1)/3} \bmod s \quad \text{or} \quad m \equiv g^{2(s-1)/3} \bmod s. \tag{2.4.50}$$

We have two cases:

II.a: If 3 does not divide $s-1$, i.e. $s \equiv 2 \bmod 3$, then the only solution to (2.4.49) is $m=1$.

II.b: If 3 does divide $s-1$, i.e. $s \equiv 1 \bmod 3$, then the solution set to (2.4.49) is:

$$\{1, g^{(s-1)/3}, g^{2(s-1)/3} : g \text{ is a primitive root mod } s\}. \tag{2.4.51}$$

Therefore all elements in $\mathbb{A}_s^3$ of the form $[n_1, g^{(s-1)/3}n_1, g^{2(s-1)/3}n_1]$, $n_1 \neq 0$, have stabilizer $\{1, g^{(s-1)/3}, g^{2(s-1)/3}\}$. Furthermore, all elements of this form lie in the following orbit:

$$\mathcal{O}_{[1,g^{(s-1)/3},g^{2(s-1)/3}]} = \{[m, mg^{(s-1)/3}, mg^{2(s-1)/3}] : m \in \mathbb{Z}_s^\times\}, \tag{2.4.52}$$

where

$$|\mathcal{O}_{[1,g^{(s-1)/3},g^{2(s-1)/3}]}| = (s-1)/3. \tag{2.4.53}$$

Indeed, since we have assumed that $n_1 \neq 0$, there are $s-1$ choices for n_1. However, since the order of elements in the 3-tuple does not matter, choosing n_1 is the same as choosing $g^{(s-1)/3}n_1$ or $g^{2(s-1)/3}n_1$. Therefore there are $(s-1)/3$ elements of this form, and they must all lie in the orbit $\mathcal{O}_{[1,g^{(s-1)/3},g^{2(s-1)/3}]}$. Using the same techniques as in proposition 2.4.1, we may now count the number of orbits (recall that γ_1 is the number of orbits of size $s-1$):

1. If $s \equiv 1 \bmod 3$, then there are $\gamma_1 + 2$ orbits:

$$|\mathbb{A}_s^3| = \gamma_1(s-1) + (s-1)/2 + (s-1)/3.$$

   Solving for γ_1 we get $\gamma_1 + 2 = (s^2 - 2s + 7)/6$.

2. If $s \equiv 2 \bmod 3$, then there are $\gamma_1 + 1$ orbits:

$$|\mathbb{A}_s^3| = \gamma_1(s-1) + (s-1)/2.$$

   Solving for γ_1 we get $\gamma_1 + 1 = (s^2 - 2s + 3)/6$.

$\square$

2.4.2 The Structure of the Orbits of $\mathbb{A}_s^d$

We now turn our attention to the more general setting, beginning with the following theorem which addresses the order of the orbits of $\mathbb{A}_s^d$ and the form of the elements in the orbits.

Theorem 2.4.3. *Let s be a prime number and let $1 < d < s$. If $\mathcal{O}$ is an orbit of $\mathbb{A}_s^d$ under the group action π_2, then there exists $c \in \mathbb{N}$ such that $c \mid d$ or $c \mid d-1$, and*

$$|\mathcal{O}| = (s-1)/c. \tag{2.4.54}$$

Furthermore, let g be a primitive root mod s and set

$$n_k^c := [n_k, g^{(s-1)/c} n_k, \ldots, g^{(c-1)(s-1)/c} n_k], \quad n_k \neq 0. \tag{2.4.55}$$

If $[n] \in \mathcal{O}$, then $[n]$ can be written in the form

$$[n] = \begin{cases} [n_1^c, n_2^c, \ldots, n_{d/c}^c] & \text{if } c \mid d, \\ [0, n_1^c, n_2^c, \ldots, n_{(d-1)/c}^c] & \text{if } c \mid d-1. \end{cases} \tag{2.4.56 c}$$

Proof. Let $m \in \mathbb{Z}_s^\times$; we determine which elements of $\mathbb{A}_s^d$ are stabilized by m based on the order of m. In particular, we will break the argument into two cases: $|m| = c > d$ and $|m| = c \leq d$. We begin with the former.

I. Assume $|m| = c > d$.

We show that no element in $\mathbb{A}_s^d$ can be stabilized by m. Let $[n] = [n_1, \ldots, n_d] \in \mathbb{A}_s^d$, $n_j \neq 0$ for all $j = 1, \ldots, d$, and suppose

$$\begin{aligned} m \cdot [n] &= [n], \\ \Longrightarrow m \cdot [n_1, \ldots, n_d] &= [n_1, \ldots, n_d], \\ \Longrightarrow [mn_1, \ldots, mn_d] &= [n_1, \ldots, n_d]. \end{aligned}$$

Therefore, $mn_1 \equiv n_j \bmod s$ for some $j \in \{1, \ldots, d\}$, and because the order of $n_1, \ldots, n_d$ does not matter, without loss of generality we have two choices:

$$mn_1 \equiv \begin{cases} n_1 \bmod s, \\ n_2 \bmod s. \end{cases} \tag{2.4.57}$$

If $mn_1 \equiv n_1 \bmod s$, then $m = 1$ and we have a contradiction to the assumption $|m| = c > d$. Therefore, $mn_1 \equiv n_2 \bmod s$ must hold. Continuing, we see that $mn_2 \equiv n_j \bmod s$ for some $j \in \{1, \ldots, d\}$. Without loss of generality, we now have three choices:

$$mn_2 \equiv \begin{cases} n_1 \bmod s, \\ n_2 \bmod s, \\ n_3 \bmod s. \end{cases} \tag{2.4.58}$$

If $mn_2 \equiv n_1 \bmod s$, then, combining this with the fact that $mn_1 \equiv n_2 \bmod s$, we see that $m^2 = 1$. However, this contradicts our initial assumption, and so is eliminated from consideration. Similarly, $mn_2 \equiv n_2 \bmod s$ implies $m = 1$ and again leads to a contradiction. Therefore, $mn_2 \equiv n_3 \bmod s$ must hold. Continuing in the same manner, we see:

$$
\begin{aligned}
mn_1 &\equiv & mn_1 &\equiv n_2 \bmod s, \\
mn_2 &\equiv & m^2 n_1 &\equiv n_3 \bmod s, \\
mn_3 &\equiv & m^3 n_1 &\equiv n_4 \bmod s, \\
& & \vdots & \\
mn_{d-1} &\equiv & m^{d-1} n_1 &\equiv n_d \bmod s.
\end{aligned}
$$

Therefore, we must have $mn_d \equiv m^d n_1 \equiv n_1 \bmod s$, which implies $m^d = 1$. Since this contradicts our initial assumption, we see that no element $m \in \mathbb{Z}_s^\times$ with $|m| = c > d$ can stabilize an element of $\mathbb{A}_s^d$ of the form $[n_1, \ldots, n_d]$, $n_j \neq 0$ for all $j = 1, \ldots, d$. The argument for elements of the form $[0, n_1, \ldots, n_{d-1}]$, $n_j \neq 0$ for all $j = 1, \ldots, d-1$, follows similarly.

II. Assume $|m| = c \leq d$.

We show an element of $\mathbb{A}_s^d$ is stabilized by m if and only if $c \mid d$ or $c \mid d-1$. First, suppose $c \nmid d$ and $c \nmid d-1$. Therefore, there exists $q, r \in \mathbb{Z}$ such that

$$d = qc + r, \quad q \geq 0, \;\; 1 < r < c. \tag{2.4.59}$$

Let $[n] = [n_1, \ldots, n_d] \in \mathbb{A}_s^d$, $n_j \neq 0$ for all $j = 1, \ldots, d$, and suppose $m \cdot [n] = [n]$.

Following the same argument as in part I of this proof, we see:

$$\begin{aligned} mn_1 &\equiv & mn_1 &\equiv n_2 \bmod s, \\ mn_2 &\equiv & m^2 n_1 &\equiv n_3 \bmod s, \\ & & \vdots & \\ mn_{c-1} &\equiv & m^{c-1} n_1 &\equiv n_c \bmod s, \\ mn_c &\equiv & m^c n_1 &\equiv n_1 \bmod s, \end{aligned}$$

where the last line results from the fact that $|m| = c \leq d$. Continuing, we see there are two possibilities for mn_{c+1}:

$$mn_{c+1} \equiv \begin{cases} n_j \bmod s \text{ for some } j \in \{1, \ldots, c\}, \\ n_{c+2} \bmod s. \end{cases} \tag{2.4.60}$$

If $mn_{c+1} \equiv n_j \bmod s$ for some $j \in \{1, \ldots, c\}$, then $mn_{c+1} \equiv mn_{j-1} \bmod s$, where $n_0 := n_c \bmod s$. However, this would imply that $n_{c+1} \equiv n_{j-1} \bmod s$, a contradiction. Therefore, $mn_{c+1} \equiv mn_{c+2} \bmod s$ must hold, and we can continue with the previous line of reasoning to obtain:

$$\begin{aligned} mn_{c+1} &\equiv & mn_{c+1} &\equiv n_{c+2} \bmod s, \\ mn_{c+2} &\equiv & m^2 n_{c+1} &\equiv n_{c+3} \bmod s, \\ & & \vdots & \\ mn_{2c-1} &\equiv & m^{c-1} n_{c+1} &\equiv n_{2c} \bmod s, \\ mn_{2c} &\equiv & m^c n_{c+1} &\equiv n_{c+1} \bmod s. \end{aligned}$$

Continuing with the pattern that has now been established, we arrive at:

$$\begin{aligned} mn_{qc+1} &\equiv & mn_{qc+1} &\equiv n_{qc+2} \bmod s, \\ mn_{qc+2} &\equiv & m^2 n_{qc+1} &\equiv n_{qc+3} \bmod s, \\ & & \vdots & \\ mn_{qc+r-1} &\equiv & m^{r-1} n_{qc+1} &\equiv n_{qc+r} \bmod s. \end{aligned}$$

We must then have:

$$mn_{qc+r} \equiv m^r n_{qc+1} \equiv n_{qc+1} \bmod s, \tag{2.4.61}$$

which in turn implies $m^r = 1$, a contradiction. Therefore, no element of $\mathbb{A}_s^d$ of the form $[n_1, \ldots, n_d]$, $n_j \neq 0$ for all $j = 1, \ldots, d$, can be stabilized by an $m \in \mathbb{Z}_s^\times$ with $|m| = c \leq d$ such that $c \nmid d$ and $c \nmid d-1$. The argument for elements of $\mathbb{A}_s^d$ of the form $[0, n_1, \ldots, n_{d-1}]$, $n_j \neq 0$ for all $j = 1, \ldots, d-1$, follows similarly.

We now shift our attention to $m \in \mathbb{Z}_s^\times$ such that $c \mid d$ or $c \mid d-1$. In either case there exists a $q \in \mathbb{Z}$ such that,

$$d = qc, \;\; q \geq 0, \quad \text{or} \quad d-1 = qc, \;\; q \geq 0. \tag{2.4.62}$$

Using the same argument that we just completed, we see that if $c \mid d$ then m stabilizes certain elements of the form $[n_1, \ldots, n_d]$, $n_j \neq 0$ for all $j = 1, \ldots, d$, whereas if $c \mid d-1$ then m stabilizes certain elements of the form $[0, n_1, \ldots, n_{d-1}]$, $n_j \neq 0$ for all $j = 1, \ldots, d-1$. The only difference in reasoning comes at the end, where in this case we do not run into a contradiction. Furthermore, looking back at the above reasoning, we see all elements $[n_1, \ldots, n_d] \in \mathbb{A}_s^d$ stabilized by m must satisfy:

$$mn_{jc+k} \equiv m^k n_{jc+1} \equiv n_{jc+k+1}, \quad \forall \; j = 0, \ldots, q-1, \;\; k = 1, \ldots, c-1, \tag{2.4.63}$$

where $d = qc$ or $d - 1 = qc$, depending on the type of element of $\mathbb{A}_s^d$. By equation (2.4.63), any element in $\mathbb{A}_s^d$ stabilized by m can be written in one of two general forms:

$$[n] = \begin{cases} [n_1, mn_1, \ldots, m^{c-1}n_1, \ldots, n_{\frac{d}{c}}, mn_{\frac{d}{c}}, \ldots, m^{c-1}n_{\frac{d}{c}}] \\ [0, n_1, mn_1, \ldots, m^{c-1}n_1, \ldots, n_{\frac{d-1}{c}}, mn_{\frac{d-1}{c}}, \ldots, m^{c-1}n_{\frac{d-1}{c}}], \end{cases} \tag{2.4.64}$$

where $n_j \neq 0$ and $n_j \neq n_k$ for all $j, k = 1, \ldots, d/c$ or $j, k = 1, \ldots, (d-1)/c$, depending on the form of $[n]$. Also, since $|m| = c$, there must exist a primitive root mod s, g, such that

$$m = g^{(s-1)/c}, \tag{2.4.65}$$

noting that $c \mid s - 1$ since the order of any group element must divide the order of the group. Combining equations (2.4.64) and (2.4.65) gives (2.4.56 c).

In order to prove (2.4.54), we exploit the fact that

$$|\mathcal{O}_{[n]}| = \frac{s-1}{|(\mathbb{Z}_s^\times)_{[n]}|}. \tag{2.4.66}$$

By (2.4.66), we need only compute the stabilizer of $[n]$ in $\mathbb{Z}_s^\times$, that is $(\mathbb{Z}_s^\times)_{[n]}$. But (2.4.64) and (2.4.65) easily give

$$(\mathbb{Z}_s^\times)_{[n]} = \{g^{l(s-1)/c} : l = 0, \ldots, c-1\}. \tag{2.4.67}$$

Clearly $|(\mathbb{Z}_s^\times)_{[n]}| = c$, thus proving (2.4.54). □

Before counting the number of orbits $\mathbb{A}_s^d$, we prove two lemmas that simplify this task. The first shows that the choice of g in (2.4.55) does not matter.

Lemma 2.4.4. *If g_1 and g_2 are two primitive roots mod s, and $n_1 \in \mathbb{Z}_s$, $n_1 \neq 0$, then*

$$[n_1, g_1^{(s-1)/c} n_1, \ldots, g_1^{(c-1)(s-1)/c} n_1] = [n_1, g_2^{(s-1)/c} n_1, \ldots, g_2^{(c-1)(s-1)/c} n_1]. \quad (2.4.68)$$

Proof. Since g_1 and g_2 are both primitive roots mod s, the sets $\{1, g_1^{(s-1)/c}, \ldots, g_1^{(c-1)(s-1)/c}\}$ and $\{1, g_2^{(s-1)/c}, \ldots, g_2^{(c-1)(s-1)/c}\}$ are both complete solution sets to $x^c \equiv 1 \bmod s$. Therefore $(n_1, g_2^{(s-1)/c} n_1, \ldots, g_2^{(c-1)(s-1)/c} n_1)$ is a rearrangement of $(n_1, g_1^{(s-1)/c} n_1, \ldots, g_1^{(c-1)(s-1)/c} n_1)$, and the lemma follows. □

The second lemma shows that the representation given by (2.4.56 c) is not unique and gives the instances where confusion can occur.

Lemma 2.4.5. *Let $[n] \in \mathbb{A}_s^d$ such that $[n]$ can be written in the form (2.4.35 b). If $c \mid b$, then $[n]$ can be written in the form (2.4.56 c) as well.*

Proof. We assume $[n] = [\tilde{n}_1^b, \tilde{n}_2^b, \ldots, \tilde{n}_{d/b}^b]$ and show that we can rewrite this as $[n] = [n_1^c, n_2^c, \ldots, n_{d/c}^c]$. If $[n] = [0, \tilde{n}_1^b, \tilde{n}_2^b, \ldots, \tilde{n}_{(d-1)/b}^b]$ then a similar proof shows how to rewrite this as $[n] = [0, n_1^c, n_2^c, \ldots, n_{(d-1)/c}^c]$. Recall

$$\tilde{n}_k^b = [\tilde{n}_k, g^{(s-1)/b} \tilde{n}_k, \ldots, g^{(b-1)(s-1)/b} \tilde{n}_k].$$

Let $a = b/c$ and set $n_1 = \tilde{n}_1$; we want to construct n_1^c out of elements of $\tilde{n}_1^b$, where:

$$\tilde{n}_1^b = [\tilde{n}_1, g^{(s-1)/b} \tilde{n}_1, \ldots, g^{(b-1)(s-1)/b} \tilde{n}_1^b].$$

Since the order of elements does not matter, we may pick them however we like and rearrange them as we wish. We have that n_1^c is formed out of the following elements of $\tilde{n}_1^b$:

$$\begin{aligned} n_1^c &= [\tilde{n}_1, g^{a(s-1)/b}\tilde{n}_1, \ldots, g^{(c-1)a(s-1)/b}\tilde{n}_1] \\ &= [n_1, g^{a(s-1)/b}n_1, \ldots, g^{(c-1)a(s-1)/b}n_1] \\ &= [n_1, g^{a(s-1)/ca}n_1, \ldots, g^{(c-1)a(s-1)/ca}n_1] \\ &= [n_1, g^{(s-1)/c}n_1, \ldots, g^{(c-1)(s-1)/c}n_1]. \end{aligned}$$

Likewise, set $n_k = \tilde{n}_k$ for $k = 2, \ldots, d/b$, and construct n_k^c in a similar manner. For the next c-tuple, set $n_{\frac{d}{b}+1} = g^{(s-1)/b}\tilde{n}_1$. We then have:

$$\begin{aligned} n_{\frac{d}{b}+1}^c &= [g^{(s-1)/b}\tilde{n}_1, g^{(a+1)(s-1)/b}\tilde{n}_1, \ldots, g^{((c-1)a+1)(s-1)/b}\tilde{n}_1] \\ &= [n_{\frac{d}{b}+1}, g^{a(s-1)/b}n_{\frac{d}{b}+1}, \ldots, g^{(c-1)a(s-1)/b}n_{\frac{d}{b}+1}] \\ &= [n_{\frac{d}{b}+1}, g^{(s-1)/c}n_{\frac{d}{b}+1}, \ldots, g^{(c-1)(s-1)/c}n_{\frac{d}{b}+1}]. \end{aligned}$$

In general,

$$n_{\frac{jd}{b}+k} = g^{j(s-1)/b}\tilde{n}_k, \quad \forall\ j = 0, \ldots, a-1, \ \ k = 1, \ldots, d/b, \tag{2.4.69}$$

and the resulting $n_{\frac{jd}{b}+k}^c$ follows similarly. □

2.4.3 Proof of Theorem 2.2.1

Using theorem 2.4.3 as well as lemmas 2.4.4 and 2.4.5, we now count the number of orbits of $\mathbb{A}_s^d$. By proposition 2.3.2 this is the same as counting the number of inequivalent harmonic frames, and so will complete the proof of theorem 2.2.1. Let γ_c denote the total number of orbits of $\mathbb{A}_s^d$ with $(s-1)/c$ elements. Then, by theorem 2.4.3, the total number of orbits of $\mathbb{A}_s^d$ is given by

$$\gamma_1 + \sum_{\substack{c|d \\ c>1}} \gamma_c + \sum_{\substack{c|d-1 \\ c>1}} \gamma_c. \tag{2.4.70}$$

Notice the similarity between equations (2.4.70) and (2.2.28). In fact, we shall prove that

$$\gamma_c = \alpha_c, \quad \forall\ c \in \mathbb{N} \text{ such that } c \mid s-1 \text{ and } c \mid d \text{ or } c \mid d-1. \tag{2.4.71}$$

Theorem 2.4.6. *Let s be a prime number, $1 < d < s$, $c \mid s-1$, $c > 1$, and let β_c denote the cumulative order of all orbits of size $(s-1)/c$. Furthermore, let γ_c denote the number of orbits of $\mathbb{A}_s^d$ of size $(s-1)/c$, so that*

$$\gamma_c = \frac{c\beta_c}{s-1}.$$

If $c \mid d$, then β_c is given by the following backwards recursive formula:

$$\beta_c = \frac{(s-1)(s-1-c)\cdots(s-1-(\frac{d}{c}-1)c)}{c^{\frac{d}{c}}(d/c)!} - \sum_{\substack{c<b<s \\ c|b,\, b|d}} \binom{s-1}{b}\gamma_b.$$

If $c \mid d-1$, then β_c is given by the following backwards recursive formula:

$$\beta_c = \frac{(s-1)(s-1-c)\cdots(s-1-(\frac{d-1}{c}-1)c)}{c^{\frac{d-1}{c}}((d-1)/c)!} - \sum_{\substack{c<b<s \\ c|b,\, b|d-1}} \binom{s-1}{b}\gamma_b.$$

The number of orbits of $\mathbb{A}_s^d$ of size $s-1$, denoted γ_1, is given by:

$$\gamma_1 = \frac{1}{s-1}\binom{s}{d} - \sum_{\substack{c|d \\ c>1}} \frac{\gamma_c}{c} - \sum_{\substack{c|d-1 \\ c>1}} \frac{\gamma_c}{c}.$$

Proof. We prove the formula for β_c when $c \mid d$, noting that the proof is identical for the case when $c \mid d-1$. In order to accomplish this task, we will build up the formula using combinatorial arguments. By theorem 2.4.3, the elements we are counting are of the form $[n_1^c, n_2^c, \ldots, n_{d/c}^c]$, where

$$n_k^c = [n_k, g^{(s-1)/c} n_k, \ldots, g^{(c-1)(s-1)/c} n_k], \quad n_k \neq 0.$$

It is clear then, that we have $s-1$ choices for n_1, $s-1-c$ choices for n_2, $s-1-2c$ choices for n_3, and so on. Continuing to the end, we see there are $s-1-(d/c-1)c$ choices for $n_{d/c}$. Furthermore, by lemma 2.4.4, the choice of g does not matter, and so does not add any new elements to count. Therefore, at the moment, we have

$$(s-1)(s-1-c)\cdots(s-1-(d/c-1)c) \tag{2.4.72}$$

elements. Fixing the choice of n_1 temporarily, it is clear that if we chose any of $g^{(s-1)/c} n_1, \ldots, g^{(c-1)(s-1)/c} n_1$ instead of n_1, then we would have a rearranged version of n_1^c. However, the order of elements does not matter in $\mathbb{A}_s^d$, and so these choices are in fact the same as choosing n_1. Since there are c such elements (including n_1), the number of distinct choices for n_1 is in fact $(s-1)/c$. Similarly, we must divide the number of choices for each n_k by a factor of c, thus giving

$$\frac{(s-1)(s-1-c)\cdots(s-1-(d/c-1)c)}{c^{d/c}} \tag{2.4.73}$$

elements. Furthermore, again recalling that the order of elements does not matter, we see that the order in which we choose $n_1, \ldots, n_{d/c}$ does not matter either. Consequently, we are now down to

$$\frac{(s-1)(s-1-c)\cdots(s-1-(d/c-1)c)}{c^{d/c}(d/c)!} \tag{2.4.74}$$

elements. We note that equation (2.4.74) gives the number of elements of the form (2.4.56 c). However, we are not counting all elements of the form (2.4.56 c), but only those that are in an orbit of size $(s-1)/c$. In fact, by lemma 2.4.5 any element in an orbit of size $(s-1)/b$, where $c \mid b$ and $b \mid d$, can be rewritten as $[n_1^c, n_2^c, \ldots, n_{d/c}^c]$. Therefore, we must subtract all elements in orbits of size $(s-1)/b$, where $c \mid b$ and $b \mid d$, thus giving:

$$\beta_c = \frac{(s-1)(s-1-c)\cdots(s-1-(\frac{d}{c}-1)c)}{c^{\frac{d}{c}}(d/c)!} - \sum_{\substack{c<b<s \\ c|b,\, b|d}} \binom{s-1}{b} \gamma_b. \tag{2.4.75}$$

The equation for γ_1 follows from

$$|\mathbb{A}_s^d| = \gamma_1(s-1) + \sum_{\substack{c|d \\ c>1}} \binom{s-1}{c} \gamma_c + \sum_{\substack{c|d-1 \\ c>1}} \binom{s-1}{c} \gamma_c, \tag{2.4.76}$$

and the fact that $|\mathbb{A}_s^d| = \binom{s}{d}$. $\square$

Example 2.4.7. We apply theorem 2.4.6 for the case when $d = 3$ and $s \equiv 1 \bmod 3$. In this case, $c = 3$ divides d as well as $s-1$, while $c = 2$ divides $d-1$ as well as $s-1$. Therefore we compute:

$$\beta_3 = \frac{s-1}{3} \quad \text{and} \quad \beta_2 = \frac{s-1}{2},$$

which in turn gives:

$$\gamma_3 = 1 \quad \text{and} \quad \gamma_2 = 1.$$

Thus,

$$\begin{aligned} \gamma_1 &= \frac{1}{s-1}\binom{s}{3} - \frac{1}{3} - \frac{1}{2} \\ &= \frac{s^2-2s}{6} - \frac{2}{6} - \frac{3}{6} \\ &= \frac{s^2-2s-5}{6}, \end{aligned}$$

and so the total number of orbits is:

$$\begin{aligned}\gamma_1+\gamma_2+\gamma_3 &= \frac{s^2-2s-5}{6}+1+1\\ &= \frac{s^2-2s+7}{6}.\end{aligned}$$

Notice this is the same result as proposition 2.4.2.

2.5 The Symmetry Group

We now turn our attention to the symmetry group of prime order harmonic frames. The following theorem proves the existence of a particular subgroup of $\mathrm{Sym}(\Phi_n)$ that is dependent on the generators $n_1,\ldots,n_d$ as well as the order of $\mathcal{O}_{[n]}$.

Theorem 2.5.1. *Let $\mathcal{O}_{[n]}$ be an orbit of $\mathbb{A}_s^d$ such that $|\mathcal{O}_{[n]}|=(s-1)/c$, and let Φ_n be the harmonic frame that corresponds to $\mathcal{O}_{[n]}$ under the one-to-one correspondence described by proposition 2.3.2. Then*

$$\langle \mathrm{diag}(\omega^{n_1},\ldots,\omega^{n_d}),Q\rangle \subseteq \mathrm{Sym}(\Phi_n),$$

where $\mathrm{diag}(\omega^{n_1},\ldots,\omega^{n_d})$ *denotes a $d\times d$ matrix with $\omega^{n_1},\ldots,\omega^{n_d}$ on the diagonal and zeros elsewhere, $\omega=e^{2\pi i/s}$, Q is a $d\times d$ permutation matrix dependent on Φ_n, and $|\langle Q\rangle|=c$.*

Proof. Similar to the proof of theorem 2.3.1, we shall almost exclusively consider Φ as a $d\times s$ matrix. We recall that $U\in\mathrm{Sym}(\Phi_n)$ if and only if there exists an $s\times s$ permutation matrix P such that

$$U\Phi=\Phi P. \tag{2.5.77}$$

First using the left hand side of (2.5.77), we have

$$(U\Phi)^\star(U\Phi) = \Phi^\star U^\star U\Phi = \Phi^\star\Phi, \tag{2.5.78}$$

and then equivalently for the right hand side of (2.5.77),

$$(\Phi P)^\star(\Phi P) = P^\star\Phi^\star\Phi P. \tag{2.5.79}$$

Combining (2.5.78) and (2.5.79) we obtain the following necessary condition for (2.5.77),

$$\Phi^\star\Phi = P^\star\Phi^\star\Phi P,$$

or equivalently,

$$P\Phi^\star\Phi P^\star = \Phi^\star\Phi. \tag{2.5.80}$$

The matrix $\Phi^\star\Phi$ is called the Gram matrix and has the following form:

$$(\Phi^\star\Phi)_{j,k} = \langle\varphi_k, \varphi_j\rangle = \sum_{l=1}^{d} e^{2\pi i n_l (k-j)/s}, \quad \forall\ j,k = 0,\dots,s-1. \tag{2.5.81}$$

Two elements $\langle\varphi_k, \varphi_j\rangle$ and $\langle\varphi_{k'}, \varphi_{j'}\rangle$ of $\Phi^\star\Phi$ are equal if and only if

$$\sum_{l=1}^{d} e^{2\pi i n_l (k-j)/s} = \sum_{l=1}^{d} e^{2\pi i n_l (k'-j')/s}. \tag{2.5.82}$$

Using the same minimum polynomial argument as the one found in the proof of theorem 2.3.1, we see that (2.5.82) holds for off diagonal elements of $\Phi^\star\Phi$ if and only if there exists a permutation $\mu \in S_d$ such that

$$n_l(k-j) \equiv n_{\mu(l)}(k'-j') \bmod s, \quad \forall\ l = 1,\dots,d,\ \ k \neq j,\ \ k' \neq j'. \tag{2.5.83}$$

(2.5.83) is in fact the same condition as (2.3.39), and so we may define the following equivalence relation between the off diagonal entries of $\Phi^\star\Phi$ and the elements of $\mathbb{A}_s^d$:

$$\langle\varphi_k, \varphi_j\rangle \sim (k-j \bmod s)\cdot[n], \quad k \neq j. \tag{2.5.84}$$

For the diagonal entries of $\Phi^\star\Phi$, we define the representative $[0]$ as

$$[0] := [\underbrace{0, \ldots, 0}_{d}], \tag{2.5.85}$$

and extend our equivalence relation to diagonal elements:

$$\langle \varphi_j, \varphi_j \rangle \sim [0]. \tag{2.5.86}$$

In order to ease notation, we set $0 \cdot [n] := [0]$, and thus can write $k \cdot [n]$ for all $k \in \mathbb{Z}_s$. Combining (2.5.84) and (2.5.86), we see $\sim$ induces an equivalence relation between the set of inner products, $\{\langle \varphi_j, \varphi_k \rangle : j, k = 0, \ldots, s-1\}$, and the set $\mathbb{A}_s^d \cup \{[0]\}$. Defining the matrix G as

$$G_{j,k} := (k - j) \cdot [n], \quad \forall\ j, k \in \mathbb{Z}_s, \tag{2.5.87}$$

we then have an equivalence relation between $\Phi^\star\Phi$ and G:

$$\Phi^\star\Phi \sim G. \tag{2.5.88}$$

Combining (2.5.80) with (2.5.88) gives the following necessary condition for (2.5.77) to hold:

$$PGP^\star = G. \tag{2.5.89}$$

Returning to (2.5.87), we see G has the form:

$$G = \begin{pmatrix} a_0 & a_{s-1} & a_{s-2} & \cdots & a_2 & a_1 \\ a_1 & a_0 & a_{s-1} & a_{s-2} & \cdots & a_2 \\ a_2 & a_1 & a_0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & a_{s-1} & a_{s-2} \\ a_{s-2} & \cdots & a_2 & a_1 & a_0 & a_{s-1} \\ a_{s-1} & a_{s-2} & \cdots & a_2 & a_1 & a_0 \end{pmatrix}, \tag{2.5.90}$$

where $a_k = k \cdot [n]$ for all $k \in \mathbb{Z}_s$. Therefore G is a circulant matrix, and is completely determined by its first column vector. The permutation matrix

$$T := \begin{pmatrix} 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \\ 1 & 0 & 0 & 0 & \cdots & 0 \end{pmatrix}, \tag{2.5.91}$$

is called the basic circulant permutation matrix. A matrix A can be written in the form

$$A = \sum_{k=0}^{s-1} a_k T^k, \tag{2.5.92}$$

if and only if A is circulant. Therefore, G can be written in the form (2.5.92), and as such, it is clear that

$$T^k G (T^k)^\star = G, \quad \forall\ k = 0, \dots, s-1. \tag{2.5.93}$$

A simple computation shows that when $U = \operatorname{diag}(\omega^{n_1}, \dots, \omega^{n_d})$, one has

$$U^k \Phi = \Phi T^k, \quad \forall\ k = 0, \dots, s-1. \tag{2.5.94}$$

Thus, regardless of the size $\mathcal{O}_{[n]}$,

$$\operatorname{diag}(\omega^{kn_1}, \dots, \omega^{kn_d}) \in \operatorname{Sym}(\Phi_n), \quad \forall\ k = 0, \dots, s-1. \tag{2.5.95}$$

Note this proves the theorem for the case $|\mathcal{O}_{[n]}| = s - 1$.

To prove the existence of the matrix $Q \in \operatorname{Sym}(\Phi_n)$ with $|\langle Q \rangle| = c$, suppose that Φ_n corresponds to $\mathcal{O}_{[n]}$ such that $|\mathcal{O}_{[n]}| = (s-1)/c$, where $c > 1$. Note that by theorem 2.4.3 we have

$$g^{k(s-1)/c} m \cdot [n] = m \cdot [n], \quad \forall\ m \in \mathbb{Z}_s^\times,\ \ k = 1, \dots, c, \tag{2.5.96}$$

and in particular

$$g^{k(s-1)/c} \cdot [n] = [n], \quad \forall\ k = 1, \ldots, c. \tag{2.5.97}$$

Therefore, the action of $g^{(s-1)/c}$ on n defines a permutation $\rho \in S_d$ such that

$$(n_{\rho^k(1)}, \ldots, n_{\rho^k(d)}) = g^{k(s-1)/c} \cdot (n_1, \ldots, n_d), \quad \forall\ k = 1, \ldots, c. \tag{2.5.98}$$

Since a permutation of the generators $n_1, \ldots, n_d$ is equivalent to a permutation of the rows of Φ, (2.5.98) implies the existence of a $d \times d$ permutation matrix Q, where Q is the matrix equivalent of ρ, as well as an $s \times s$ permutation matrix P_0, such that

$$Q^k \Phi = \Phi P_0^k, \quad \forall\ k = 1, \ldots, c. \tag{2.5.99}$$

In other words, $Q \in \mathrm{Sym}(\Phi_n)$, and since $\mathrm{diag}(\omega^{n_1}, \ldots, \omega^{n_d}) \in \mathrm{Sym}(\Phi_n)$ as well, we must have

$$\langle \mathrm{diag}(\omega^{n_1}, \ldots, \omega^{n_d}), Q \rangle \subseteq \mathrm{Sym}(\Phi_n). \tag{2.5.100}$$

$\square$

Corollary 2.5.2. *Let $\mathcal{O}_{[n]}$ be an orbit of $\mathbb{A}_s^d$ such that $|\mathcal{O}_{[n]}| = s - 1$, and let Φ_n be the harmonic frame that corresponds to $\mathcal{O}_{[n]}$ under the one-to-one correspondence described by proposition 2.3.2. Then*

$$\mathrm{Sym}(\Phi_n) = \langle \mathrm{diag}(\omega^{n_1}, \ldots, \omega^{n_d}) \rangle,$$

where $\mathrm{diag}(\omega^{n_1}, \ldots, \omega^{n_d})$ *denotes a $d \times d$ matrix with $\omega^{n_1}, \ldots, \omega^{n_d}$ on the diagonal and zeros elsewhere, and $\omega = e^{2\pi i/s}$.*

Proof. Recall the matrices G and T from the proof of theorem 2.5.1, as given by equations (2.5.87) and (2.5.91), respectively. We will show that $P = T^k$, $k =$

$0, \ldots, s-1$, are the only matrices satisfying the necessary condition given by equation (2.5.89). Combining the fact that $\mathcal{O}_{[n]} = \{m \cdot [n] : m \in \mathbb{Z}_s^\times\}$ with the assumption that $|\mathcal{O}_{[n]}| = s-1$, we have

$$k \cdot [n] = k' \cdot [n] \iff k \equiv k' \bmod s. \tag{2.5.101}$$

Furthermore, let $\sigma \in S_s$ be the permutation corresponding to the permutation matrix P. Equation (2.5.89) can be rewritten as

$$(\sigma(j) - \sigma(k)) \cdot [n] = (j-k) \cdot [n], \quad \forall\ j, k \in \mathbb{Z}_s. \tag{2.5.102}$$

Combining equations (2.5.101) and (2.5.102), one obtains

$$\sigma(j) - \sigma(k) = j - k, \quad \forall\ j, k \in \mathbb{Z}_s. \tag{2.5.103}$$

One can think of (2.5.103) as a system of s^2 linear equations in the s variables $\sigma(0), \ldots, \sigma(s-1)$, with the two added constraints:

1. $\sigma(k) \in \mathbb{Z}_s$ for all $k \in \mathbb{Z}_s$,

2. $\sigma(j) = \sigma(k)$ if and only if $j = k$.

Clearly (2.5.103) is an overdetermined system. However, (2.5.103) has $s-1$ independent equations, given by:

$$\begin{aligned}
\sigma(1) - \sigma(0) &\equiv 1 \bmod s \\
\sigma(2) - \sigma(0) &\equiv 2 \bmod s \\
&\vdots \\
\sigma(s-1) - \sigma(0) &\equiv s-1 \bmod s.
\end{aligned}$$

Thus $\sigma(0)$ is a free variable, and can be assigned any value from $\mathbb{Z}_s$. The remaining values of σ are then given by:

$$\sigma(j) \equiv j + \sigma(0) \bmod s, \quad \forall\ j = 1, \ldots, s-1. \tag{2.5.104}$$

In conclusion, there are s possible permutations, each corresponding to a different value of $\sigma(0)$. In particular, we have the following correspondence:

$$\sigma(0) = k \iff P = T^k. \tag{2.5.105}$$

$\square$

The following conjecture asserts that the subgroup described in theorem 2.5.1 in fact is the symmetry group for all prime order harmonic frames, not just those corresponding to orbits of size $s-1$.

Conjecture 2.5.3. *Let $\mathcal{O}_{[n]}$ be an orbit of $\mathbb{A}_s^d$ such that $|\mathcal{O}_{[n]}| = (s-1)/c$, and let Φ_n be the harmonic frame that corresponds to $\mathcal{O}_{[n]}$ under the one-to-one correspondence described by proposition 2.3.2. Then*

$$\mathrm{Sym}(\Phi_n) = \langle \mathrm{diag}(\omega^{n_1}, \ldots, \omega^{n_d}), Q\rangle,$$

where $\mathrm{diag}(\omega^{n_1}, \ldots, \omega^{n_d})$ denotes a $d \times d$ matrix with $\omega^{n_1}, \ldots, \omega^{n_d}$ on the diagonal and zeros elsewhere, $\omega = e^{2\pi i/s}$, Q is a $d \times d$ permutation matrix dependent on Φ_n, and $|\langle Q\rangle| = c$.

2.6 Closing remarks

We have enumerated all harmonic frames for $\mathbb{C}^d$ with s elements, where s is a prime number. A natural question is how to extend these results to all s. Certain problems arise, however, with the techniques used in this chapter, since in several instances the fact that s is prime is a key element. In particular, for a general s, distinct harmonic frames will arise from groups other than $\mathbb{Z}_s$. Also, even for those harmonic frames that do come from $\mathbb{Z}_s$, new representations must be developed since in general $\mathbb{Z}_s^\times \subseteq \{1, \ldots, s\}$.

Chapter 3

Frame Based Kernel Methods

3.1 Introduction to Multispectral and Hyperspectral Imagery Data

When a camera takes a picture, reflected light from the subject is passed through three filters: red, green, and blue. The resulting bands are then combined to form a color image; see figure 3.1. Multispectral and hyperspectral cameras,

Figure 3.1: Color image decomposition

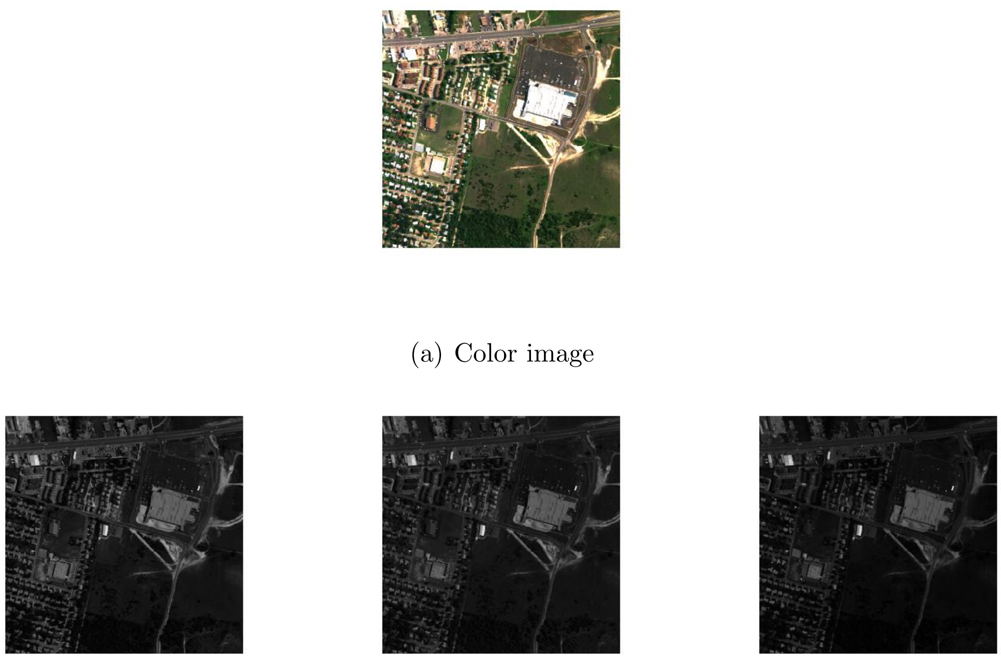

(a) Color image

(b) Red band

(c) Green band

(d) Blue band

on the other hand, are in a sense a generalization of a regular camera. Rather than filter the reflected light through red, green, and blue filters, these cameras are able to measure reflectance at a multitude of different wavelengths; see figure 3.2. Observing the sample bands in figure 3.2, one notices that the reflectance is measured at wavelengths far beyond the visible spectrum (recall the visible spectrum is about 380 – 750 nm). Unlike standard cameras, the purpose of multispectral and hyperspectral cameras is not to create a color image, but rather to collect as much information about a particular scene as possible. This additional information can then be used for many different tasks, including (see [30] for more details):

- target detection
- material mapping
- material identification
- mapping details of surface properties

What figure 3.2 does not illustrate, though, is the central difference between multispectral imagery (MSI) data sets and hyperspectral imagery (HSI) data sets. While the bands are spread across the spectrum, there are large gaps where no measurements are shown. HSI data sets are in fact characterized by the narrowness and contiguous nature of their measurements, leaving few if any large spectral gaps. One can imagine stacking each of the bands and forming a cube, as illustrated by figure 3.3. Given this abundance of spectral information, HSI data sets are generally spectrally overdetermined, and are thus able to distinguish between spectrally similar materials. In order to achieve the contiguous nature of the measurements,

Figure 3.2: Selected bands of a hyperspectral data set

(a) 412 nm

(b) 468 nm

(c) 543 nm

(d) 650 nm

(e) 808 nm

(f) 1014 nm

(g) 1237 nm

(h) 1451 nm

(i) 1648 nm

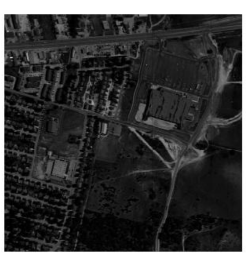

(j) 1992 nm

(k) 2144 nm

(l) 2284 nm

Figure 3.3: Hyperspectral data cube

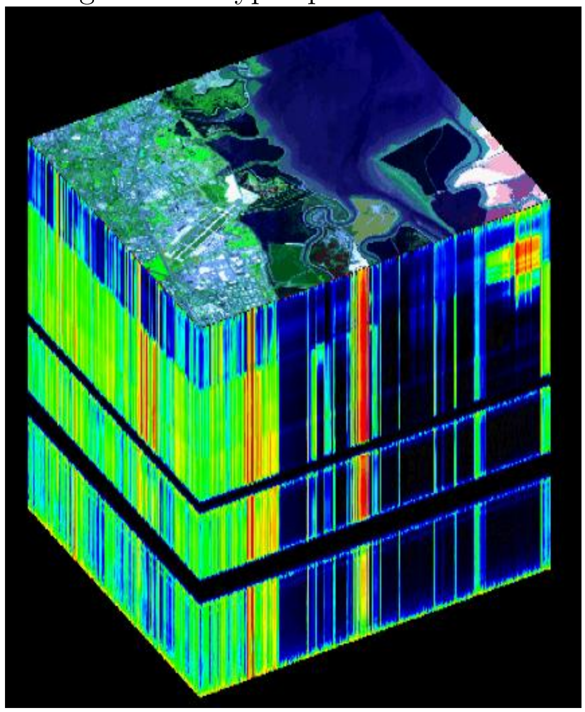

usually there are hundreds of spectral bands. MSI data sets on the other hand have anywhere from 4 spectral bands up to one hundred. In practical applications, the number of pixels is on the order of hundreds of thousands, sometimes even millions.

Mathematically speaking, we can model a MSI/HSI data set in the following way. Let X denote a MSI/HSI data cube with dimensions $N_1 \times N_2 \times D$, where N_1 and N_2 are the spatial dimensions (length and width), and D is the spectral dimension. Thus we have $N = N_1 N_2$ pixels, each measured at D different wavelengths. Most of the time it will be easier to think of X as a list, and so we set $X = \{x_i : i = 1, \dots, N\} \subset \mathbb{R}^D$, where each x_i corresponds to a pixel.

We will use MSI/HSI data for the purpose of material classification: given a list of potential classes within a data set, we aim to correctly classify each pixel as a certain class. Since HSI data sets are spectrally overdetermined, there are usually less classes than the number of wavelengths measured. The MSI data sets that we examine will also have this property, although in general this is not true. A traditional method for classification in HSI data is through the use of endmember extraction algorithms. *Endmembers* are defined as a collection of the scene's constituent spectra. If $E = \{e_i : i = 1, \ldots, s\}$ are endmembers for the HSI data set X, then the *linear mixture model* is

$$x_i = \sum_{j=1}^{s} \alpha_{i,j} e_j + N_{x_i}, \quad \forall\, x_i \in X, \tag{3.1.1}$$

where N_{x_i} is a noise vector. The set $\{\alpha_{i,j} : i = 1, \ldots, N,\ j = 1, \ldots, s\}$ are the coefficients, and it is usually assumed that they satisfy the following two conditions:

$$\alpha_{i,j} \geq 0, \quad \forall\, i = 1, \ldots, N,\ j = 1, \ldots, s, \tag{3.1.2}$$

$$\sum_{j=1}^{s} \alpha_{i,j} = 1, \quad \forall\, i = 1, \ldots, N. \tag{3.1.3}$$

Let $\alpha_{i,\cdot} = (\alpha_{i,1}, \ldots, \alpha_{i,s})$ and let $\tilde{\alpha} = (\tilde{\alpha}_1, \ldots, \tilde{\alpha}_s) \in \mathbb{R}^s$. Two common endmember coefficient sets are given by:

$$\alpha_{i,\cdot} = \arg\min_{\tilde{\alpha}} \|x_i - \sum_{j=1}^{s} \tilde{\alpha}_j e_j\|_{\ell^2} \quad \text{subject to} \quad (3.1.2), (3.1.3), \tag{3.1.4}$$

$$\alpha_{i,\cdot} = \arg\min_{\tilde{\alpha}} \|x_i - \sum_{j=1}^{s} \tilde{\alpha}_j e_j\|_{\ell^2} + \tau_i \|\tilde{\alpha}\|_{\ell^1} \quad \text{subject to} \quad (3.1.2), (3.1.3), \tag{3.1.5}$$

where τ_i is a positive real number. Most endmember extraction algorithms determine E as a subset of X, i.e. it is assumed that the endmembers lie within the given data set. There are several endmember extraction algorithms, including N-FINDR [35], ORASIS [11], Pixel Purity Index [10], and Support Vector Data Description (SVDD) [6, 31]; see also [15, 21].

3.2 Overview of New Algorithm

The main objective of this part of the thesis is to introduce a new algorithm for the purposes of material classification in MSI and HSI data sets. This algorithm is based on the theory of frames and dimension reduction, in particular kernel eigenmap methods.

As stated above, traditional endmember algorithms determine a subset $E \subset X$ by which to represent the elements of X. Another way to view this is that they are determining a low dimensional subspace of interest, in this case $\text{span}(E)$. The algorithm presented here uses techniques from dimension reduction to give an alternate method for determining a low dimensional space of interest. We shall use kernel eigenmap methods to map the high dimensional space X to a low dimensional space Y. Unlike endmember algorithms, Y will not be determined as a subspace X, but rather through a nonlinear mapping.

We will then construct frame by which to represent the space Y. Akin to endmember extraction algorithms, this frame Φ can be a subset of Y. We will also present methodologies by which to construct a data dependent frame from scratch that is not a subset of Y. Regardless of how they are constructed, unlike endmember sets, these frames will provide overcomplete representations, a fact we shall exploit.

There are many techniques for dimension reduction, e.g., Principal Component Analysis (PCA) [23], Locally Linear Embedding (LLE) [28], Isomap [32], genetic algorithms, and neural networks. We are interested in a subfamily of these techniques known as kernel eigenmap methods. These include Kernel PCA [29], LLE, Hessian

LLE (HLLE) [18], and Laplacian eigenmaps [7]. Kernel eigenmap methods require two steps.

1. Construction of a symmetric, positive semi-definite kernel (a matrix), K, for given data and a specific type of dimension reduction problem to solve.

2. Diagonalization of K to obtain the eigenmaps (eigenvectors).

We shall interpret the data and kernel dependent Hilbert space $\mathbb{K}$, mentioned in section 3.4 for general kernel eigenmap methods, in terms of the theory of frames. Frames provide non-orthogonal overcomplete signal decompositions. In dealing with dimension reduction, our experiments to compare spectral signatures illustrate that different classes are almost never orthogonal, whereas eigenmap methods provide processed orthogonal decompositions. On the contrary, frame elements are not necessarily orthogonal. As such, for given data, they can be constructed to reflect empirical non-orthogonal angular relations between classes. Driven by this mathematical modeling in terms of frames, and inspired solely by given data, we describe an innovative methodology to achieve class separability and object identification.

3.3 Kernel Eigenmap Methods

Given a high dimensional data set $X = \{x_i : i = 1, \ldots, N\} \subset \mathbb{R}^D$, we assume that the data points x_i in fact lie on a low dimensional manifold M^d, where d is the dimension of the manifold, and $d < D$. As an example, see figure 3.4, where we have a collection of points in $\mathbb{R}^2$ that lie on a one dimensional manifold. Dimension reduction methods construct a mapping from $\mathbb{R}^D$ to $\mathbb{R}^d$, and in particular map X

Figure 3.4: Points in $\mathbb{R}^2$ on a one dimensional manifold

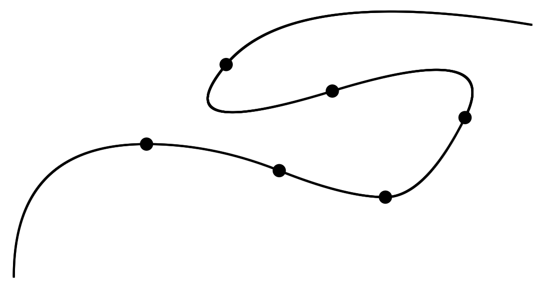

to low dimensional coordinates $Y = \{y_i : i = 1, \ldots, N\} \subset \mathbb{R}^d$, where $x_i \mapsto y_i$. The main goal of these methods is to have the new coordinates Y preserve the underlying geometric structure of the manifold M^d.

Kernel eigenmap methods are a subset of dimension of reduction methods. The key component of these methods is the construction of a data dependent, $N \times N$, symmetric, positive semi-definite kernel K:

$$K_{i,j} = K(x_i, x_j), \quad \forall\, i, j = 1, \ldots, N. \tag{3.3.6}$$

The kernel K is then diagonalized, and the d significant eigenvectors of K are retained. Let $v_1, \ldots, v_d \in \mathbb{R}^N$ denote these eigenvectors; the new low dimensional coordinates Y are then given by:

$$y_i = (v_1(i), \ldots, v_d(i)), \quad \forall\, i = 1, \ldots, N. \tag{3.3.7}$$

3.3.1 Spectral Clustering

We now give a more in depth overview of the theory behind kernel eigenmap methods. Recall

$$X = \{x_i : i = 1, \ldots, N\} \subset \mathbb{R}^D. \tag{3.3.8}$$

We then define a distance $\rho : \mathbb{R}^D \times \mathbb{R}^D \longrightarrow \mathbb{R}^+$ such that

$$\begin{aligned} \rho(x_i, x_j) &= \rho(x_j, x_i), \quad \forall\, i, j = 1, \ldots, N, \\ \rho(x_i, x_i) &= 0, \quad \forall\, i = 1, \ldots, N. \end{aligned} \tag{3.3.9}$$

Since $\mathbb{R}^D$ is a vector space, we can define ρ in terms of a norm $\|\cdot\|$, where $\rho(x_i, x_j) = \|x_i - x_j\|$, $\|\cdot\| : \mathbb{R}^D \longrightarrow \mathbb{R}^+$, and $\|0\| = 0$. We then compute

$$\begin{aligned} \widetilde{A} &= (\widetilde{A}_{i,j})_{i,j=1}^N, \\ \widetilde{A}_{i,j} &= \rho(x_i, x_j). \end{aligned} \tag{3.3.10}$$

$\widetilde{A}$ has many nonzero entries and is therefore computationally intensive to diagonalize. We think of $\widetilde{A}$ as global information, since it gives the 'distance' between any two points in X.

For each $x_i \in X$, let

$$\mathcal{N}_k(x_i) = \{k \text{ nearest neighbors of } x_i \text{ with respect to } \rho\}. \tag{3.3.11}$$

In order to find $\mathcal{N}_k(x_i)$ for a fixed i, we order the elements of $\{\widetilde{A}_{i,j}\}_{j=1}^N$:

$$0 = \widetilde{A}_{i,i} \leq \widetilde{A}_{i,\sigma(1)} \leq \widetilde{A}_{i,\sigma(2)} \leq \ldots \leq \widetilde{A}_{i,\sigma(k)} \leq \ldots \leq \widetilde{A}_{i,\sigma(N-1)}, \tag{3.3.12}$$

where $\sigma \in S_N$. We then set

$$\mathcal{N}_k(x_i) = \{x_{\sigma(j)} : j = 1, \ldots, k\}. \tag{3.3.13}$$

Our adjacency matrix $A = (A_{i,j})$ is then given by

$$A_{i,j} = \begin{cases} 1, & x_j \in \mathcal{N}_k(x_i), \\ 0, & \text{otherwise.} \end{cases} \tag{3.3.14}$$

A has zeros down its diagonal and is not necessarily symmetric. If we want A to be symmetric we could use ε-balls instead of the k nearest neighbors, in which case we would replace $\mathcal{N}_k(x_i)$ with $B_\varepsilon(x_i) = \{x_j : \rho(x_i, x_j) < \varepsilon\}$.

We now define our directed graph $G = \{X, E\}$ where

$$E \subset X \times X \text{ and } (x_i, x_j) \in U \iff A_{i,j} = 1. \tag{3.3.15}$$

The weight matrix $\widetilde{W} = (\widetilde{W}_{i,j})$ is then given by

$$\widetilde{W}_{i,j} = h\left(\widetilde{A}^2_{i,j}\right) \cdot A_{i,j} \tag{3.3.16}$$

where h has exponential decay at ∞, e.g. $h(x) = e^{-x}$. Finally, we define a normalizing matrix $D = (D_{i,j})$ where

$$D_{i,j} = \begin{cases} \sum_l \widetilde{W}_{i,l}, & i = j, \\ 0, & i \neq j. \end{cases} \tag{3.3.17}$$

We then set

$$W = D^{-1}\widetilde{W}. \tag{3.3.18}$$

W contains local information on the relative distances between points; we can think of $W_{i,j}$ as the probability of walking from x_i to x_j. We now examine the following diagram, which shows two points, x_i and x_j, that are *two* edges apart.

$$\begin{tikzcd}
x_i \arrow[r, "W_{i,p}"] \arrow[dr, "W_{i,q}"'] & x_p \arrow[r, "W_{p,j}"] & x_j \\
& x_q \arrow[ur, "W_{q,j}"'] &
\end{tikzcd}$$

We see that the probability of walking from x_i to x_j is given by $W_{i,p}W_{p,j} + W_{i,q}W_{q,j}$. But this is just an example of the following identity

$$W^2_{i,j} = \sum_{\substack{x_p \in \mathcal{N}_k(x_i) \\ x_j \in \mathcal{N}_k(x_p)}} W_{i,p}W_{p,j}. \tag{3.3.19}$$

Thus $W_{i,j}^2$ is the probability of walking from x_i to x_j in exactly two steps. More generally we have

$$W_{i,j}^l = \text{ the probability of walking from } x_i \text{ to } x_j \text{ in exactly } l \text{ steps.} \qquad (3.3.20)$$

We now look at the following example. Consider the following graph, depicted in figure 3.5, where we assume the probability of walking to any given neighbor is equal to that of some other neighbor (the arrows illustrate this point). Furthermore,

Figure 3.5: Diffusion distance

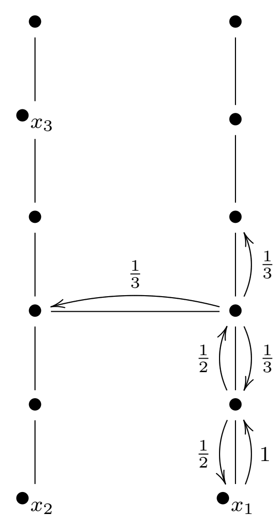

assume the graph is embedded in $\mathbb{R}^2$ and take ρ as the Euclidean distance. Clearly $\widetilde{A}_{1,2} < \widetilde{A}_{1,3}$. However, $W_{1,2}^5 = W_{1,3}^5$. This is called diffusion similarity, and thus we call W the diffusion matrix.

Recall we have a directed graph $G = \{X, E\}$, where $X = \{x_i : i = 1, \ldots, N\}$. Let $f : X \to \mathbb{R}$. Since X has N elements, we can think of f as a vector in $\mathbb{R}^N$. However, each $x_i \in \mathbb{R}^D$, so f takes $\mathbb{R}^D$ into $\mathbb{R}$, i.e. $f : \mathbb{R}^D \to \mathbb{R}$. We want to define

$\nabla f \in \mathbb{R}^D$.

Think of the directed edges of G as vectors. Define

$$u_j = \lambda x_j + (1-\lambda)x_i, \quad 0 \leq \lambda \leq 1. \tag{3.3.21}$$

Assuming f is a linear function on $\mathbb{R}^D \to \mathbb{R}$ we have $f(u_j) = \lambda f(x_j) + (1-\lambda)f(x_i)$. Thus we set

$$\nabla_{u_j} f = f(x_j) - f(x_i). \tag{3.3.22}$$

We now want to reparameterize to take into account the weights

$$\begin{aligned} &0 \leq \lambda \leq \tfrac{1}{\sqrt{W_{i,j}}}, \\ &u_j = \sqrt{W_{i,j}}\lambda x_j + \sqrt{W_{i,j}}(1-\lambda)x_i, \\ &\nabla_{u_j} f = \sqrt{W_{i,j}}(f(x_j) - f(x_i)). \end{aligned} \tag{3.3.23}$$

We now define

$$\operatorname{div} f = \sum_{j=1}^{N} \nabla_{u_j} f, \tag{3.3.24}$$

$$\Delta = -\operatorname{div} \cdot \nabla. \tag{3.3.25}$$

Therefore

$$\begin{aligned} \Delta f(x_i) &= \sum_{x_j \in \mathcal{N}_k(x_i)} W_{i,j}(f(x_j) - f(x_i)) \\ &= -f(x_i) + \sum_{x_j \in \mathcal{N}_k(x_i)} W_{i,j} f(x_j) \qquad (3.3.26) \\ &= (-I + W)f. \qquad (3.3.27) \end{aligned}$$

(3.3.26) shows that our definition of Δ is in fact a good one, i.e., $\Delta f(x_i) = 0$ if and only if f satisfies the mean value property on $\mathcal{N}_k(x_i)$. (3.3.27) shows that $W - I$ is the discrete Laplacian operator.

We now describe the two types of kernels used in this research.

3.3.2 Locally Linear Embedding

For each $x_i \in X$ compute k-nearest neighbors of x_i, $\mathcal{N}_k(x_i)$, and construct the directed graph described in section 3.3.1. Furthermore, we assume that the graph G is connected. The weights are computed by solving the following minimization problem:

$$W = \arg\min_{\widetilde{W}} \sum_{i=1}^{N} \left| x_i - \sum_{x_j \in \mathcal{N}(x_i)} \widetilde{W}_{i,j} x_j \right|^2, \tag{3.3.28}$$

subject to the constraint:

$$\sum_{j=1}^{N} \widetilde{W}_{i,j} = 1, \quad \forall\, i = 1, \ldots, N. \tag{3.3.29}$$

Notice that equation (3.3.28) can be rewritten line by line:

$$W_{i,\cdot} = \arg\min_{\widetilde{W}_{i,\cdot}} \left| x_i - \sum_{x_j \in \mathcal{N}(x_i)} \widetilde{W}_{i,j} x_j \right|^2. \tag{3.3.30}$$

The LLE kernel is then defined as:

$$K = (I - W)^{\star}(I - W), \tag{3.3.31}$$

which, when compared to (3.3.27), one sees that

$$K = \Delta^2. \tag{3.3.32}$$

The eigenvectors of K will have nonnegative eigenvalues; there will be one with an eigenvalue of zero. The d significant eigenvectors are given by those eigenvectors that correspond the d smallest nonzero eigenvalues.

3.3.3 Laplacian Eigenmaps

Like LLE, we use k-nearest neighbors to determine the neighborhoods of each $x_i \in X$ and we construct the graph G. The weights are defined as:

$$\widetilde{W}_{i,j} = \begin{cases} \exp\{-\|x_i - x_j\|_{\ell^2}^2/\sigma\}, & \text{if } x_j \in \mathcal{N}(x_i) \text{ or } x_i \in \mathcal{N}(x_j) \\ 0, & \text{otherwise,} \end{cases} \tag{3.3.33}$$

where σ is a positive real number. Define the $N \times N$, diagonal matrix D that same as in equation (3.3.17):

$$D_{i,j} = \begin{cases} \sum_{l=1}^{N} \widetilde{W}_{i,l}, & i = j \\ 0, & i \neq j \end{cases} \tag{3.3.34}$$

We then set the Laplacian eigenmap kernel to be:

$$K = D - \widetilde{W}. \tag{3.3.35}$$

Notice that

$$K = -D\Delta. \tag{3.3.36}$$

For Laplacian eigenmaps, the eigenmaps are obtained by solving the following generalized eigenvector problem:

$$Kf = \lambda Df. \tag{3.3.37}$$

Like LLE, the we select the d eigenvectors corresponding to the d smallest non-zero eigenvalues.

3.4 Theoretical Foundations of the Algorithm

This algorithm, which appears mathematical and is mathematically sound, is the foundation for our computational work. It is not, however, a direct transcription of the actual computations, but rather the inspiration for them. The differences between the theoretical work here and the actual computations are detailed in section 3.5.

Given a data set $X = \{x_i : i = 1, \ldots, N\} \subset \mathbb{R}^D$, we create the kernel, K, using existing kernel methods such as locally linear embedding (LLE), Laplacian eigenmaps, and Hessian LLE. $K \in \mathcal{M}_{N,N}(\mathbb{R})$ is a square matrix of size N and rank r. Furthermore, we construct K to be positive semi–definite, i.e., for all vectors $f \in \mathbb{R}^N$, we have $f^\star K f \geq 0$, where $f^\star$ is the transpose of the complex conjugate of f.

As with any square matrix, we can diagonalize K. In particular, there exists $V \in \mathcal{M}_{N,N}(\mathbb{R})$ and $\{\lambda_i : i = 1, \ldots, N\} \subset \mathbb{R}^+$ such that

$$VV^\star = V^\star V = I, \tag{3.4.38}$$

and

$$K = V\mathrm{diag}(\lambda_i)V^\star. \tag{3.4.39}$$

Furthermore, by Mercer's theorem [26], there exists a reproducing kernel Hilbert space $\mathbb{K}$ and a set of vectors $\{\psi_i : i = 1, \ldots, N\} \subset \mathbb{K}$, such that for all $i, j = 1, \ldots, N$,

$$\langle \psi_i, \psi_j \rangle_{\mathbb{K}} = K_{i,j}. \tag{3.4.40}$$

We note that $\dim\mathbb{K} = \mathrm{rank}(K) = r$. Furthermore, let $\mathcal{B} = \{b_i : i = 1, \ldots, r\}$ be any orthonormal basis for $\mathbb{K}$. Then there exists $\widetilde{Y} \in \mathcal{M}_{N,r}(\mathbb{R})$ such that for each $i = 1, \ldots, N$ we have

$$\psi_i = \sum_{j=1}^{r} \widetilde{Y}_{i,j} b_j, \tag{3.4.41}$$

where

$$\widetilde{Y}_{i,j} = \langle \psi_i, b_j \rangle_{\mathbb{K}}. \tag{3.4.42}$$

It is clear then that

$$\widetilde{Y}\widetilde{Y}^{\star} = K, \tag{3.4.43}$$

and that $\text{rank}(\widetilde{Y}) = \text{dim}\mathbb{K} = \text{rank}(K) = r$. We note that this $N \times r$ matrix $\widetilde{Y}$, while not the same as the set $Y = \{y_i : i = 1, \dots, N\} \subset \mathbb{R}^d$ first introduced in section 3.3, plays a similar role.

Lemma 3.4.1. *We can choose $\mathcal{B}$ so that*

$$\widetilde{Y} = V'_{N,r}\text{diag}(\sqrt{\lambda_i})_{r,r}, \tag{3.4.44}$$

where $V'_{N,r}$ is the matrix of columns of V (i.e. the eigenvectors of K) that correspond to non-zero λ_i.

Proof. Consider the singular value decomposition of $\widetilde{Y}$: there exists $U_1 \in \mathcal{M}_{N,N}(\mathbb{R})$, $U_2 \in \mathcal{M}_{r,r}(\mathbb{R})$, and $\{\omega_i : i = 1, \dots, N\} \subset \mathbb{R}$ such that

$$U_1 U_1^{\star} = U_1^{\star} U_1 = I, \tag{3.4.45}$$

$$U_2 U_2^{\star} = U_2^{\star} U_2 = I, \tag{3.4.46}$$

and

$$\widetilde{Y} = U_1 \Omega U_2^{\star}, \tag{3.4.47}$$

where $\Omega = \left[\dfrac{\text{diag}_{r,r}(\omega_i)}{\mathbf{0}}\right]_{N,r}$. Note that $\mathbf{0}$ denotes a block of zeros. Hence,

$$K = \widetilde{Y}\widetilde{Y}^{\star} = U_1 \Omega U_2^{\star} U_2 \Omega^{\star} U_1^{\star} = U_1 \Omega \Omega^{\star} U_1^{\star} \tag{3.4.48}$$

where

$$\Omega\Omega^\star = \left[\frac{\operatorname{diag}(|\omega_i|^2)|\mathbf{0}}{\mathbf{0}|\mathbf{0}}\right]. \tag{3.4.49}$$

Therefore, the ω_i are uniquely determined (up to a phase factor) by the eigendecomposition of K, and so are the columns of U_1 that correspond to non-zero ω_i. $\square$

We pick $s \in \mathbb{Z}$, with $r \leq s \leq N$. Let $\Phi = \{\varphi_i : i = 1, \ldots, s\}$ be a FUNTF for $\mathbb{K}$. There exists a coefficient matrix $C \in \mathcal{M}_{N,s}(\mathbb{R})$ such that for each $i = 1, \ldots, N$ we have

$$\psi_i = \sum_{j=1}^{s} C_{i,j}\varphi_j. \tag{3.4.50}$$

We can choose

$$C_{i,j} = \langle \psi_i, \varphi_j \rangle_{\mathbb{K}} \tag{3.4.51}$$

but these are not the unique coefficients for which equation (3.4.51) is valid. We can also represent Φ in terms of the basis $\mathcal{B}$, i.e., there exists $Z \in \mathcal{M}_{s,r}(\mathbb{R})$ such that for each $i = 1, \ldots, s$, we have

$$\varphi_i = \sum_{j=1}^{r} Z_{i,j} b_j. \tag{3.4.52}$$

We note that Φ is a FUNTF if and only if $Z^\star Z = \frac{s}{r}I$, or alternatively, if and only if $\|Z^\star Z\|_{FRO} = \frac{s^2}{r}$. Combining equations (3.4.42) and (3.4.52), we have

$$\langle \psi_i, \varphi_j \rangle_{\mathbb{K}} = \sum_{l=1}^{r} \langle \psi_i, b_l \rangle_{\mathbb{K}} \langle b_l, \varphi_j \rangle_{\mathbb{K}} = \sum_{l=1}^{r} \widetilde{Y}_{i,l} Z_{j,l}. \tag{3.4.53}$$

Thus, one possible coefficient matrix is given by:

$$C = \widetilde{Y} Z^\star. \tag{3.4.54}$$

Inspired by the above calculations, we define how to find the FUNTF Ψ and set up a more general method for computing C. Let c_1 and c_2 denote cost functions on the space of matrices $\mathcal{M}_{N,s}(\mathbb{R})$. We want to find a basis $\mathcal{B} = \{b_i : i = 1, \ldots, r\}$ for $\mathbb{K}$ and a FUNTF $\Phi = \{\varphi_i : i = 1, \ldots, s\}$ for $\mathbb{K}$ that satisfy the following minimization problem:

$$\min_{\widetilde{\mathcal{B}}, \widetilde{\Phi}} c_1(\widetilde{Y} Z^\star), \quad \text{subject to} \quad (3.4.44), \tag{3.4.55}$$

where $\widetilde{\mathcal{B}}$ is any basis for $\mathbb{K}$ and $\widetilde{\Phi}$ is any FUNTF for $\mathbb{K}$. Recall that by equation (3.4.41) $\widetilde{Y}$ is completely determined by $\widetilde{\mathcal{B}}$ and that by equation (3.4.52) Z is completely determined by $\widetilde{\mathcal{B}}$ and $\widetilde{\Phi}$.

Given a FUNTF Φ, we then want to find a coefficient matrix C that satisfies (3.4.50). Using our second cost function c_2, we find C by solving the following minimization problem:

$$C = \arg\min_{\widetilde{C}} c_2(\widetilde{C}), \quad \text{subject to} \quad (3.4.50), \tag{3.4.56}$$

where $\widetilde{C}$ is any possible coefficient set.

3.5 The Algorithm in Practice

In practice the algorithm consists of the following five steps:

1. Landmarking

2. Kernel eigenmap method

3. Out of sample extension

4. Frame construction

5. Frame coefficients

The main differences between the actual algorithm and the theoretical ideas are the following. Steps one and three, landmarking and out of sample extension, are employed for certain kernels on large data sets so that the algorithm can run in a reasonable amount of time. Also, the frame construction and handling of kernels are viewed from a more practical point of view, and are implemented correspondingly. We give detailed explanations below.

3.5.1 Landmarking

For certain kernels to be used on large scale data sets, landmarking must be employed. For our purposes, we use landmarking only when applying the LLE kernel. Laplacian eigenmaps, with its simple kernel construction, we have found feasible on the data sets of interest. In the case of Laplacian eigenmaps, one my think of the algorithm as skipping steps one and three.

The idea of landmarking is to determine a subset of X that will be used to compute the kernel K, as opposed to using all of X. We denote this subset as

$$X_{sam} := \{x_{i_j} : j = 1, \ldots, n\} \subset X, \tag{3.5.57}$$

where n is the number of samples and we assume that $n \ll N$. To obtain X_{sam}, we sample X uniformly at random without replacement.

3.5.2 Kernel Eigenmap Methods

We apply the LLE and Laplacian eigenmap kernel eigenmap methods. LLE is used for HSI terrain data, while Laplacian eigenmaps is applied to MSI biological data. Unlike in section 3.4 where the work is done in the kernel space $\mathbb{K}$, we are forced to diagonalize K in practice and use the traditional reduced coordinates $Y = \{y_i = (v_1(i), \ldots, v_d(i)) : i = 1, \ldots, N\}$, which were described in section 3.3. As a matter of notation, we shall denote the reduced dimensional coordinates of X_{sam} as $y_{i_j} \in \mathbb{R}^d$. The reason for returning to the traditional methodology is the following: both LLE and Laplacian eigenmaps were designed to be run on a connected graph G, and as such, the rank of these kernels is $r = N - 1$. Thus $\dim\mathbb{K} = N - 1$, where in practice N is on the order of 10^6. To work in such a space is computationally infeasible, and so a subspace must be determined. The natural first subspace to try is the one given by the eigendecomposition. It should be noted that in the future the development of a true frame based kernel, with low rank, would at least theoretically be optimal.

3.5.3 Out of Sample Extension

Given the n low dimensional coordinates $\{y_{i_j} : j = 1, \ldots, n\}$ corresponding to the sampled set $X_{sam} = \{x_{i_j} : j = 1, \ldots, n\} \subset X$, we wish to extend these new coordinates to all of X via an out of sample extension [9]. To do so we extend the definition of k-nearest neighbors to include reference points $x_i \in X$ that are out of sample, i.e., for all $x_i \in X \backslash X_{sam}$, we define:

$$\mathcal{N}'_k(x_i) := \{x_{i_j} \in X_{sam} : x_{i_j} \text{ is one of the } k \text{ nearest neighbors of } x_i \text{ with respect to } \rho\}. \tag{3.5.58}$$

Notice that while the reference point may now come from X, the neighbors are still selected only from the sampled subset X_{sam}. In the case of LLE, we must similarly define weights for $x_i \in X \backslash X_{sam}$. Let $W'(x_i, \cdot) = (W'(x_i, x_{i_1}), \ldots, W'(x_i, x_{i_n}))$, and define it as:

$$W'(x_i, \cdot) = \arg \min_{\widetilde{W}'(x_i, \cdot)} \left| x_i - \sum_{x_{i_j} \in \mathcal{N}'(x_i)} \widetilde{W}'(x_i, x_{i_j}) x_{i_j} \right|^2, \tag{3.5.59}$$

subject to the constraint:

$$\sum_{j=1}^{n} \widetilde{W}'(x_i, x_{i_j}) = 1, \quad \forall\, x_i \in X \backslash X_{sam}. \tag{3.5.60}$$

The low dimensional coordinates for $x_i \in X \backslash X_{sam}$ are then given by:

$$y_i = \sum_{j=1}^{n} y_{i_j} W'(x_i, x_{i_j}), \tag{3.5.61}$$

thus giving a complete set of low dimensional coordinates $Y = \{y_i : i = 1, \ldots, N\}$ that correspond to $x_i \mapsto y_i$.

3.5.4 Frame Construction

Given that we have already departed from the theoretical methodology outlined in section 3.4, it is only natural that our frame construction algorithms deviate as well. Again we forgo working in the kernel space $\mathbb{K}$ for the simpler reduced coordinates given by $Y = \{y_i : i = 1, \ldots, N\} \subset \mathbb{R}^d$. Given the new coordinates Y, we construct a frame $\Phi = \{\varphi_i : i = 1, \ldots, s\} \subset \mathbb{R}^d$, where $s \geq d$, for the space span(Y). We then represent the coordinates Y in terms of the frame Φ. In this research we have used two methods to construct frames, detailed below.

3.5.4.1 Endmember Frames

When dealing with HSI terrain data, which is initially processed using the LLE kernel eigenmap method, we have applied existing endmember algorithms to the low dimensional coordinates Y. In particular, we have extensively tested the support vector data description (SVDD) endmember algorithm [6, 31] within this framework. The benefit of using an existing endmember algorithm such as SVDD is that it is fast, immediately available, and gives a means by which to compare our new framework with an existing endmember algorithm.

3.5.4.2 Maximum Separation Frames

We have also developed frame construction algorithms based on modified versions of the frame potential. Through the use of penalty terms, we are able to guide the frame to separate out various features within the data. More specifically, for a FUNTF $\Phi = \{\varphi_i : i = 1, \ldots, s\}$ and coordinates $Y = \{y_i : i = 1, \ldots, N\}$, define the following penalty function, p, as follows:

$$p(\varphi_i) := \sum_{j=1}^{N} |\langle y_j, \varphi_i \rangle|. \tag{3.5.62}$$

We also set:

$$p(\Phi) := \sum_{i=1}^{s} p(\varphi_i). \tag{3.5.63}$$

For a given t, $0 \leq t \leq s$, and ε, $0 \leq \varepsilon \leq 1$, we then compute a 'separated' FUNTF Φ by solving the following modified frame potential:

$$\Phi = \arg \min_{\widetilde{\Phi} \in \mathbb{S}^{d-1} \times \ldots \times \mathbb{S}^{d-1}} FP(\widetilde{\Phi})$$

$$\text{subject to} \tag{3.5.64}$$

$$\sum_{i=t+1}^{s} \frac{p(\tilde{\varphi}_i)}{p(\Phi)} = \sum_{i=t+1}^{s} \sum_{j=1}^{N} \frac{|\langle y_j, \tilde{\varphi}_i \rangle|}{p(\Phi)} < \varepsilon.$$

The above frame construction has been applied to biological data, which is first processed using Laplacian eigenmaps.

3.5.5 Frame Coefficients

Given reduced coordinates $Y = \{y_i : i = 1, \ldots, N\} \subset \mathbb{R}^d$ and a frame $\Phi = \{\varphi_i : i = 1, \ldots, s\} \subset \mathbb{R}^d$ for $\operatorname{span}(Y)$, we are left to compute coefficients $C = \{C_{i,j} : i = 1, \ldots, N;\ j = 1, \ldots, s\}$ to represent Y in terms of Φ. We use two types of coefficients: canonical and sparse.

3.5.5.1 Canonical Coefficients

Given a frame Φ, recall the definition of the frame operator S:

$$S(f) = \sum_{i=1}^{s} \langle f, \varphi_i \rangle \varphi_i. \tag{3.5.65}$$

The canonical coefficients of Ψ are then given by:

$$C_{i,j} = \langle y_i, S^{-1}(\varphi_j) \rangle, \quad \forall\, i = 1, \ldots, N,\ \ j = 1, \ldots, s. \tag{3.5.66}$$

It is well known that the canonical coefficients satisfy the following reconstruction formula:

$$y_i = \sum_{j=1}^{s} C_{i,j} \varphi_j, \quad \forall\, i = 1, \ldots, N. \tag{3.5.67}$$

Furthermore, the canonical coefficients are easy and fast to compute, especially so when Φ is a FUNTF and $S = \frac{s}{d} I$.

3.5.5.2 Sparse Coefficients

To compute sparse coefficients for a frame Φ we solve an ℓ^1 minimization problem for each $y_i \in Y$. Let $C_{i,\cdot} = (C_{i,1}, \ldots, C_{i,s})$; we then compute:

$$C_{i,\cdot} = \arg\min_{\widetilde{C}_{i,\cdot}} \|\widetilde{C}_{i,\cdot}\|_{\ell^1}, \quad \text{subject to} \quad \sum_{j=1}^{s} \widetilde{C}_{i,j}\varphi_j = y_i. \tag{3.5.68}$$

We use ℓ^1 minimization as a substitute for the following ℓ^0 minimization problem:

$$\arg\min_{\widetilde{C}_{i,\cdot}} \|\widetilde{C}_{i,\cdot}\|_{\ell^0}, \quad \text{subject to} \quad \sum_{j=1}^{s} \widetilde{C}_{i,j}\varphi_j = y_i, \tag{3.5.69}$$

where $\|f\| = \#\text{supp}(f)$. Solving (3.5.69) is NP hard and requires and exhaustive combinatorial search, thus making it intractable. Via the theory of compressed sensing [12, 13, 17], it has been shown that (3.5.68) can, in certain situations, be used as a direct substitute to (3.5.69), or at the very least, a good approximation. Furthermore, (3.5.68) is a convex optimization problem, and can be solved (reasonably quickly) using linear programming techniques. While the sparse coefficients are more computationally intensive than the canonical coefficients, they can provide enhanced separation of classes, especially when considered from a visual perspective.

Chapter 4

Empirical Results

In this chapter we present empirical results derived from running our algorithm on real world hyperspectral and multispectral data sets. These results are broken into two main categories:

1. Hyperspectral terrain data
   (a) Urban
   (b) Smith Island
2. Multispectral eye data

We give more details on each category, as well as the results, in the sections below.

4.1 Hyperspectral Terrain Data

We have two hyperspectral terrain data sets, Urban and Smith Island [3, 2, 4, 1, 5]. For these data sets we use the LLE kernel for the dimension reduction, and the SVDD endmember algorithm to construct a frame; we use both canonical and sparse coefficients. For both Urban and Smith Island we have a small subset of ground truth training data, which gives sample pixels of each class contained within the data. Since we have some ground truth, numerically based classification and comparison is possible; we also present frame coefficient images as well as class maps.

4.1.1 Classification Methodology

Once again denote our HSI data set as $X = \{x_i : i = 1, \ldots, N\} \subset \mathbb{R}^D$. Let $T \subset X$ denote our ground truth, and suppose there are q classes within T. Let $T_i \subset T$, $i = 1, \ldots, q$, denote the set of pixels corresponding to class i, so that:

$$\begin{aligned} &\textstyle\bigcup_{i=1}^{q} T_i = T, \\ &T_i \cap T_j = \emptyset, \quad \forall\, i \neq j. \end{aligned} \tag{4.1.1}$$

In order to perform classification on the set X, we first construct average representative vectors for each class T_i. Denote the elements of T_i (and therefore T as well) as:

$$T_i = \{t_{i,j} : j = 1, \ldots, q_i\}, \tag{4.1.2}$$

where q_i denotes the number of pixels in the class T_i; note that, by definition,

$$\sum_{i=1}^{q} q_i = \#T. \tag{4.1.3}$$

The average representative vector for the class T_i, denoted $\tilde{t}_i$, is given by:

$$\tilde{t}_i = \frac{1}{q_i} \sum_{j=1}^{q_i} t_{i,j}. \tag{4.1.4}$$

Let $\widetilde{T}$ denote the set of average representative vectors, i.e.,

$$\widetilde{T} = \{\tilde{t}_i : i = 1, \ldots, q\}. \tag{4.1.5}$$

We then classify each vector $x_i \in X$ by comparing x_i with the elements of $\widetilde{T}$. We use the spectral angle between x_i and each $\tilde{t}_j$ as the determining factor, where the angle between two vectors is given by:

$$\theta_{x_i,\tilde{t}_j} = \cos^{-1}\left(\frac{\langle x_i, \tilde{t}_j\rangle}{\|x_i\|\|\tilde{t}_j\|}\right). \tag{4.1.6}$$

If the angle between x_i and $\tilde{t}_{j_0}$ is smaller than the angle between x_i and all other $\tilde{t}_j$, then we place x_i in class j_0. Mathematically speaking, if

$$j_0 = \arg\min_{j=1,\dots,q} \theta_{x_i,\tilde{t}_j}, \tag{4.1.7}$$

then x_i is placed in class j_0. To obtain numerical statistics, we use the ground truth data set T, and its subsets corresponding to classes, T_i, $i = 1, \dots, q$. For each $t_{i,j} \in T_i$, we see if the spectral angle classifier indeed places $t_{i,j}$ in class i. This allows us to determine a percentage correct for each class, as well as for the ground truth data set as a whole.

Note that we can extend this classification method to reduced coordinates Y and coefficient coordinates C by simply computing everything in terms of these coordinates. In these situations, the indexes of T would remain the same, but we would now assume that $T \subset Y$ or $T \subset \{C_{i,\cdot} : i = 1, \dots, N\}$, respectively. $\widetilde{T}$ would thus be computed again in terms of these new coordinates as well. In the sections to follow we present spectral angle classification statistics based on ground truth for the original data set $X = \{x_i : i = 1, \dots, N\} \subset \mathbb{R}^D$, the LLE low dimensional coordinates $Y = \{y_i : i = 1, \dots, N\} \subset \mathbb{R}^d$, SVDD endmember coefficients $\{\alpha_{i,\cdot} : i = 1, \dots, N\}$ computed according to (3.1.4) and (3.1.5), as well as our frame coefficients $\{C_{i,\cdot} : i = 1, \dots, N\}$ computed according to (3.5.66) and (3.5.68).

4.1.2 Overview of the Trials

We have run our algorithm on the Urban and Smith Island data sets, which entails the following steps. First process the data set through the LLE kernel eigenmap method to obtain reduced coordinates Y. We then compute a frame for span(Y) using the SVDD algorithm. A frame coefficient cube C is then produced, and we run the spectral angle classification method on these new frame coefficient vectors to obtain class maps and statistical data.

There are two trials for the Urban data set and one for the Smith data set. For every trial there were three variable parameters: the reduced dimension of the LLE coordinates, d, the number of frame elements, s, and type of frame coefficients - canonical or sparse. We have run each trial through a variety of choices for d and s and computed the canonical or minimum ℓ^2 error coefficients for each iteration, depending on whether $s \geq d$ (canonical) or $s < d$ (minimum ℓ^2 error). We then ran spectral angle classification on each coefficient cube. For each trial, we highlight a particular d and s that had one of the highest overall percentages correct. For this particular d and s we then also compute the sparse coefficients, and present the corresponding statistics and maps.

There are also three different competing results, each serving as means of comparison to our algorithm. First among these is that we classify the raw, unprocessed, Urban and Smith Island data sets. Secondly, we use the LLE reduced coordinates alone, varying the reduced dimension d and selecting the best one according to percentage. Note that the range for d is the same regardless of whether it is for our algorithm or the LLE coordinates alone. We also run solely SVDD on the Urban data set, this time varying the number of endmembers, s. For each s we compute the minimum ℓ^2 error coefficients, and highlight the particular s value with the highest overall percentage. For this particular s we then also compute the mixed ℓ^2-ℓ^1 coefficients, and present these results as well. Again note, we use the same parameters when running SVDD alone as in the context of our algorithm, although in this case that does not necessarily mean that the s values are the same since there is no way to directly control the number of endmembers/frame elements that SVDD returns.

4.1.3 Urban

4.1.3.1 Description of the Urban Data Set

The Urban data set is a hyperspectral imagery data set that is freely available at:

$$http://www.agc.army.mil/Hypercube/index.html$$

The dimensions of Urban are $307 \times 307 \times 161$: that is $307 * 307 = 94249$ pixels and 161 spectral bands. A pseudocolor image of the Urban data set is given in figure 4.1. There are 932 ground truth pixels, broken into 22 distinct classes; these classes are:

1. AsphaltDrk
2. AsphaltLgt
3. Concrete01

4. VegPasture
5. VegGrass
6. VegTrees01
7. Soil01
8. Soil02
9. Soil03Drk
10. Roof01Wal
11. Roof02A
12. Roof02BGvl
13. Roof03LgtGray
14. Roof04DrkBrn
15. Roof05AChurch
16. Roof06School
17. Roof07Bright
18. Roof08BlueGrn
19. TennisCrt
20. PoolWater

21. ShadedVeg

22. ShadedPav

Figure 4.1: Pseudocolor image of Urban

4.1.3.2 Urban Trial 1

The results of Urban trial 1 were obtained with the following settings:

- Data set: X = Urban
- Kernel: LLE
- Number of neighbors: $k = 20$
- Number of samples: $\#X_{sam} = 20000$ pixels
- Frame construction: SVDD

The classification results for varying s and d and the canonical coefficients are displayed in figure 4.2. Note that not every combination of d and s have a result - due to the nature of the SVDD algorithm, it is only possible to 'guide' the number of endmembers by tweaking certain parameters, there is no direct way to select s. Also, the black line represents the line where $d = s$. We highlight the following

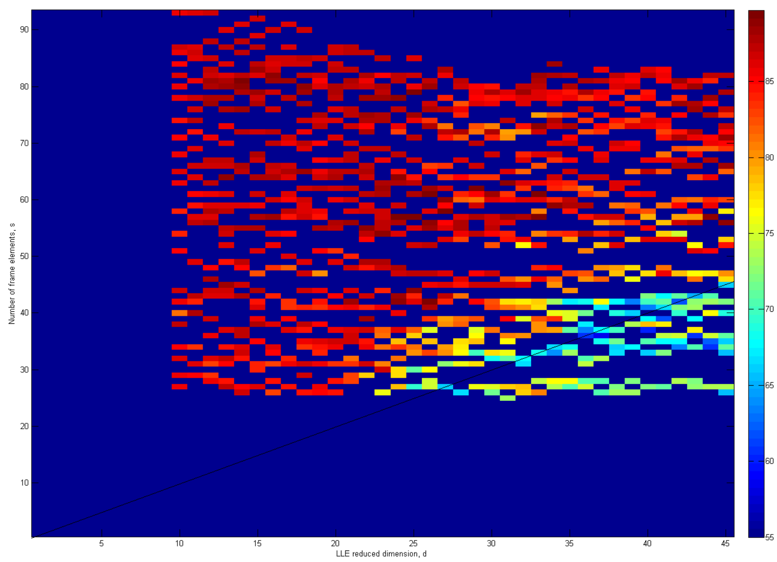

Figure 4.2: Urban trial 1 canonical coefficients classification results for varying s and d

particular cases.

Urban Trial 1 A

- Number of reduced dimensions: $d = 25$
- Number of frame elements: $s = 57$
- Type of coefficients: canonical

Statistical results for Urban trial 1 A can be found in table 4.1. Figure 4.4 shows the class map for this trial, while figures 4.8 and 4.9 show the individual class maps. Figures 4.16, 4.17, 4.18, and 4.19 show the coefficient maps for each of the frame elements.

Urban Trial 1 B

- Number of reduced dimensions: $d = 25$
- Number of frame elements: $s = 57$
- Type of coefficients: sparse

Statistical results for Urban trial 1 B can be found in table 4.2. Figure 4.5 shows the class map for this trial, while figures 4.10 and 4.11 show the individual class maps. Figures 4.20, 4.21, 4.22, and 4.23 show the coefficient maps for each of the frame elements.

4.1.3.3 Urban Trial 2

For Urban trial 2 we increased the number of neighbors over trial 1, but otherwise kept the settings the same:

- Data set: $X =$ Urban
- Kernel: LLE
- Number of neighbors: $k = 40$
- Number of samples: $\#X_{sam} = 20000$ pixels

Table 4.1: Urban trial 1 A ground truth results

	#	# correct	% correct	# false positives	# false negatives
AsphaltDrk	45	45	100%	8	0
AsphaltLgt	26	21	81%	9	5
Concrete01	64	54	84%	0	10
VegPasture	116	116	100%	3	0
VegGrass	65	63	97%	12	2
VegTrees01	123	85	69%	8	38
Soil01	52	51	98%	0	1
Soil02	24	20	83%	6	4
Soil03Drk	27	27	100%	0	0
Roof01Wal	57	57	100%	1	0
Roof02A	44	43	98%	3	1
Roof02BGvl	17	15	88%	5	2
Roof03LgtGray	12	10	83%	0	2
Roof04DrkBrn	39	39	100%	5	0
Roof05AChurch	38	34	89%	0	4
Roof06School	28	28	100%	0	0
Roof07Bright	35	35	100%	0	0
Roof08BlueGrn	21	15	71%	0	6
TennisCrt	47	42	89%	4	5
PoolWater	5	3	60%	0	2
ShadedVeg	17	9	53%	31	8
ShadedPav	30	24	80%	1	6
Total	932	836	90%	96	96

Table 4.2: Urban trial 1 B ground truth results

	#	# correct	% correct	# false positives	# false negatives
AsphaltDrk	45	45	100%	8	0
AsphaltLgt	26	21	81%	5	5
Concrete01	64	60	94%	3	4
VegPasture	116	116	100%	2	0
VegGrass	65	63	97%	10	2
VegTrees01	123	80	65%	7	43
Soil01	52	44	85%	1	8
Soil02	24	19	79%	3	5
Soil03Drk	27	26	96%	0	1
Roof01Wal	57	57	100%	1	0
Roof02A	44	43	98%	3	1
Roof02BGvl	17	15	88%	11	2
Roof03LgtGray	12	11	92%	3	1
Roof04DrkBrn	39	31	79%	4	8
Roof05AChurch	38	35	92%	0	3
Roof06School	28	28	100%	0	0
Roof07Bright	35	35	100%	0	0
Roof08BlueGrn	21	15	71%	0	6
TennisCrt	47	40	85%	3	7
PoolWater	5	3	60%	2	2
ShadedVeg	17	11	64%	38	6
ShadedPav	30	23	77%	7	7
Total	932	821	88%	111	111

- Frame construction: SVDD

The classification results for varying s and d and the canonical coefficients are displayed in figure 4.3. We highlight the following particular cases.

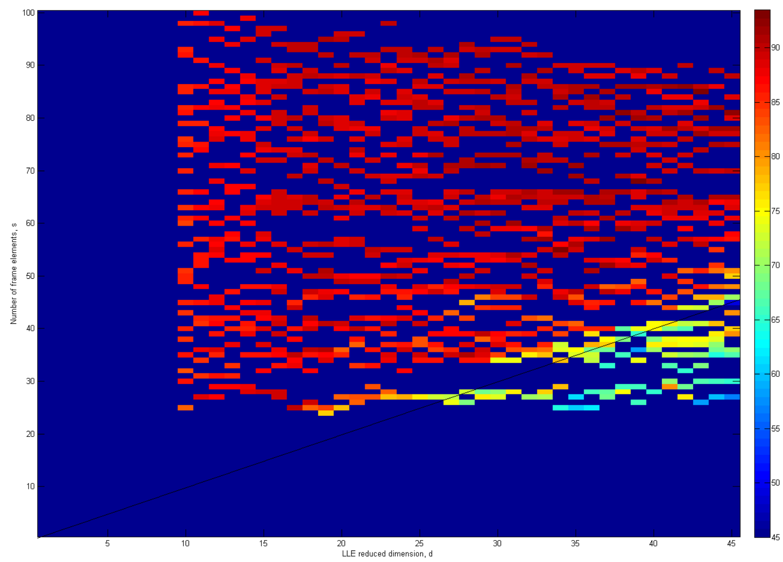

Figure 4.3: Urban trial 2 canonical coefficients classification results for varying s and d

Urban Trial 2 A

- Number of reduced dimensions: $d = 44$
- Number of frame elements: $s = 86$
- Type of coefficients: canonical

Statistical results for Urban trial 2 A can be found in table 4.3. Figure 4.6 shows the class map for this trial, while figures 4.12 and 4.13 show the individual class maps. Figures 4.24, 4.25, 4.26, 4.27, and 4.28 show the coefficient maps for each of

Table 4.3: Urban trial 2 A ground truth results

	#	# correct	% correct	# false positives	# false negatives
AsphaltDrk	45	45	100%	0	0
AsphaltLgt	26	20	77%	1	6
Concrete01	64	64	100%	1	0
VegPasture	116	116	100%	2	0
VegGrass	65	64	98%	9	1
VegTrees01	123	88	72%	4	35
Soil01	52	52	100%	0	0
Soil02	24	22	92%	1	2
Soil03Drk	27	27	100%	0	0
Roof01Wal	57	57	100%	3	0
Roof02A	44	44	100%	0	0
Roof02BGvl	17	17	100%	1	0
Roof03LgtGray	12	11	92%	1	1
Roof04DrkBrn	39	39	100%	3	0
Roof05AChurch	38	37	97%	0	1
Roof06School	28	28	100%	0	0
Roof07Bright	35	35	100%	0	0
Roof08BlueGrn	21	21	100%	0	6
TennisCrt	47	43	91%	1	4
PoolWater	5	3	60%	0	2
ShadedVeg	17	13	76%	34	4
ShadedPav	30	25	83%	0	5
Total	932	871	93%	61	61

the frame elements.

Urban Trial 2 B

- Number of reduced dimensions: $d = 44$
- Number of frame elements: $s = 86$
- Type of coefficients: sparse

Statistical results for Urban trial 2 B can be found in table 4.4. Figure 4.7 shows the class map for this trial, while figures 4.14 and 4.15 show the individual class maps. Figures 4.29, 4.30, 4.31, 4.32, and 4.33 show the coefficient maps for each of the frame elements.

Table 4.4: Urban trial 2 B ground truth results

	#	# correct	% correct	# false positives	# false negatives
AsphaltDrk	45	45	100%	1	0
AsphaltLgt	26	20	77%	0	6
Concrete01	64	61	95%	1	3
VegPasture	116	116	100%	2	0
VegGrass	65	64	98%	10	1
VegTrees01	123	92	75%	5	31
Soil01	52	52	100%	0	0
Soil02	24	24	100%	1	0
Soil03Drk	27	27	100%	0	0
Roof01Wal	57	56	98%	1	1
Roof02A	44	44	100%	3	0
Roof02BGvl	17	17	100%	0	0
Roof03LgtGray	12	12	100%	4	0
Roof04DrkBrn	39	39	100%	2	0
Roof05AChurch	38	38	100%	0	0
Roof06School	28	28	100%	0	0
Roof07Bright	35	35	100%	0	0
Roof08BlueGrn	21	20	95%	0	1
TennisCrt	47	41	87%	1	6
PoolWater	5	3	60%	0	2
ShadedVeg	17	12	71%	30	5
ShadedPav	30	25	83%	0	5
Total	932	871	93%	61	61

4.1.3.4 Urban Competing Results

Table 4.5 contains the overall results of the competing Urban results. We note that the LLE and SVDD results were obtained at the following points:

- LLE only (trial 1): $d = 45$
- LLE only (trial 2): $d = 27$
- SVDD (both coefficient cubes): $s = 8$

Table 4.5: Urban competing overall results

	#	# correct	% correct	# false pos/neg
Raw data	932	785	84%	147
LLE only (trial 1)	932	835	90%	97
LLE only (trial 2)	932	873	94%	59
SVDD only (min ℓ^2 error coeffs)	932	861	92%	71
SVDD only (mixed ℓ^2-ℓ^1 coeffs)	932	334	36%	598

4.1.3.5 Urban Class Maps

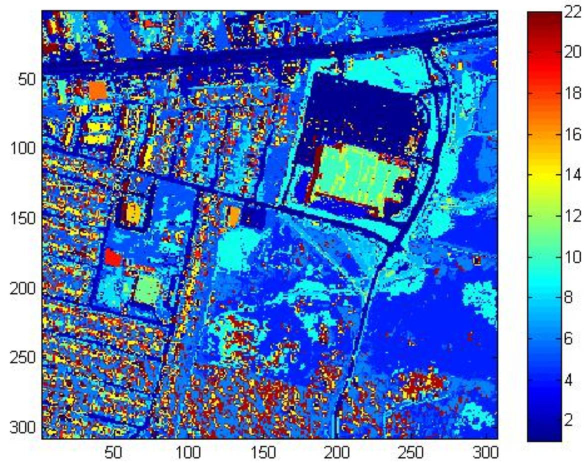

Figure 4.4: Urban trial 1 A class map

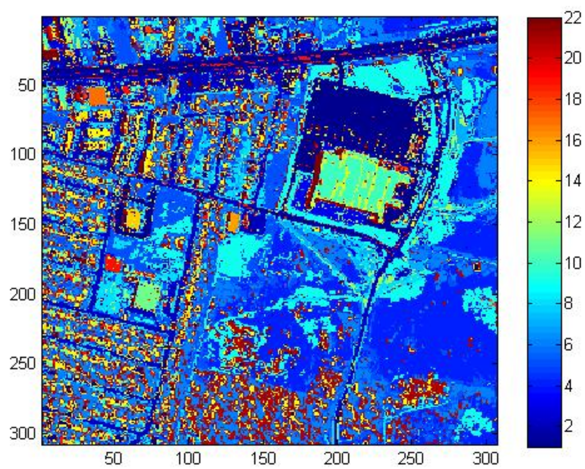

Figure 4.5: Urban trial 1 B class map

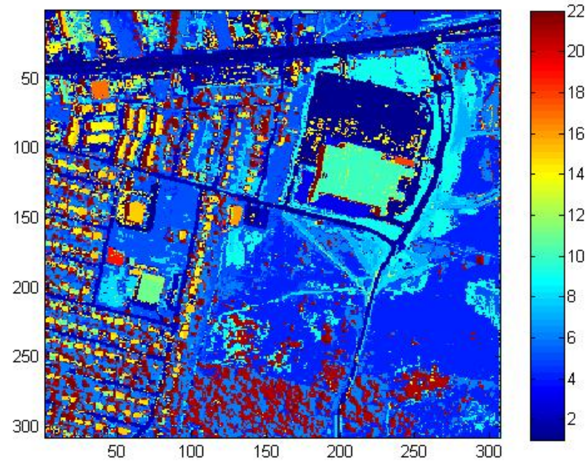

Figure 4.6: Urban trial 2 A class map

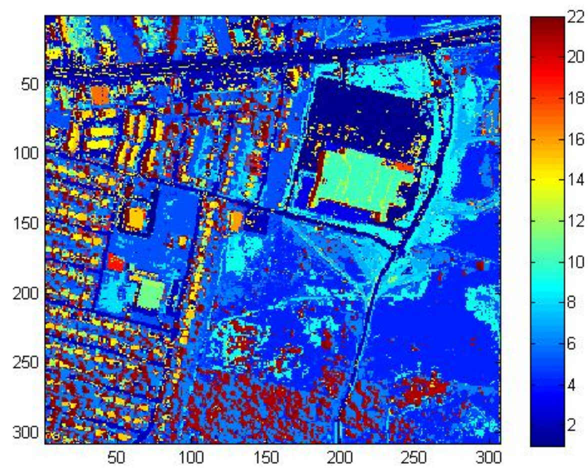

Figure 4.7: Urban trial 2 B class map

4.1.3.6 Urban Individual Class Maps

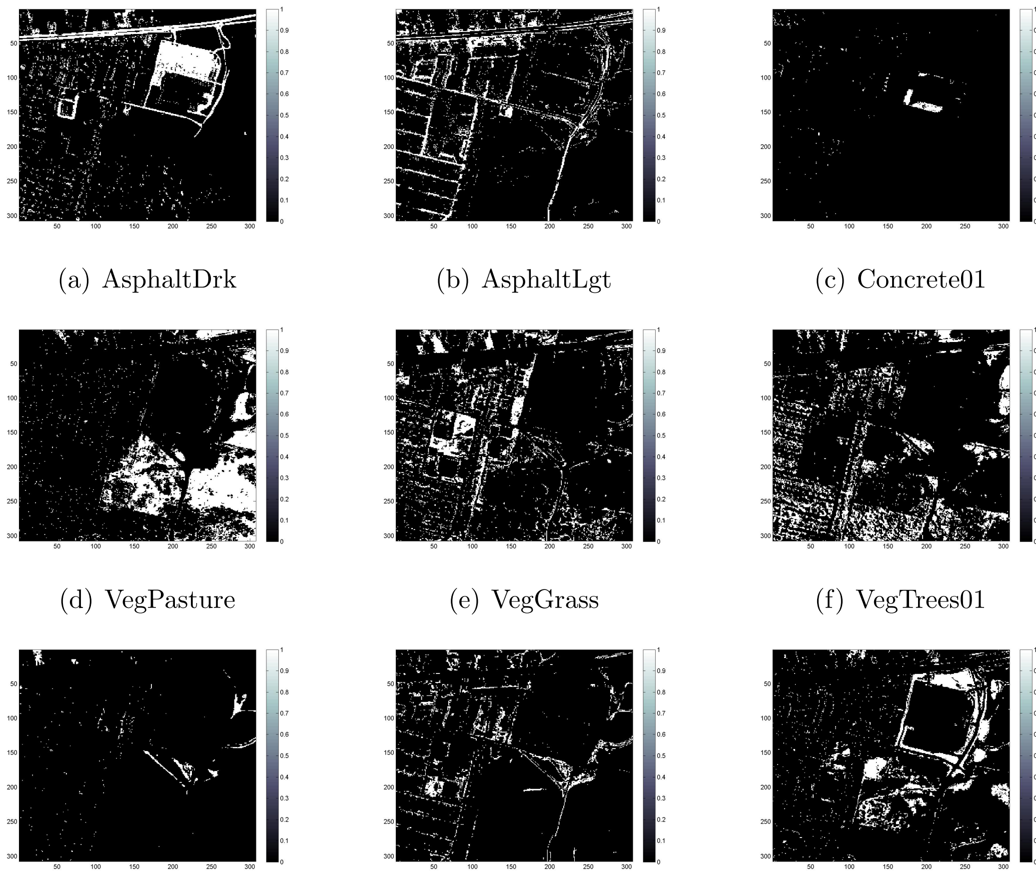

(a) AsphaltDrk (b) AsphaltLgt (c) Concrete01

(d) VegPasture (e) VegGrass (f) VegTrees01

(g) Soil01 (h) Soil02 (i) Soil03Drk

Figure 4.8: Urban trial 1 A individual class maps 1–9

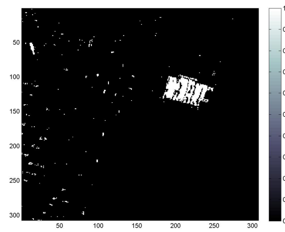

(a) Roof01Wal

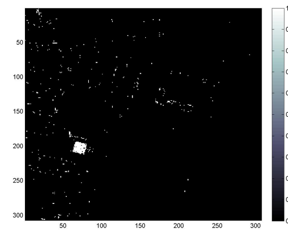

(b) Roof02A

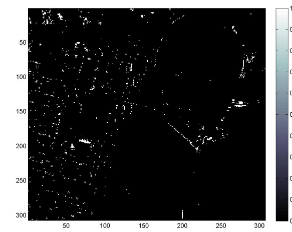

(c) Roof02BGvl

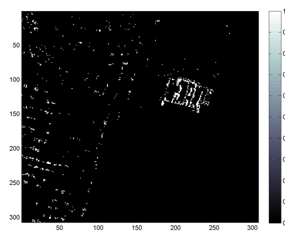

(d) Roof03LgtGray

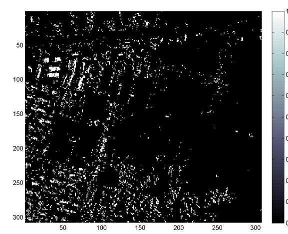

(e) Roof04DrkBrn

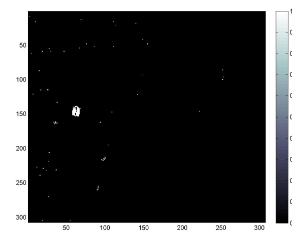

(f) Roof05AChurch

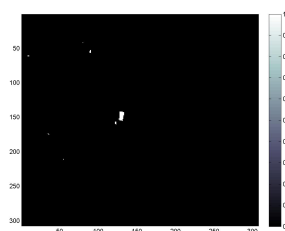

(g) Roof06School

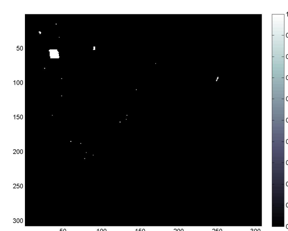

(h) Roof07Bright

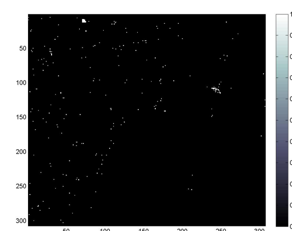

(i) Roof08BlueGrn

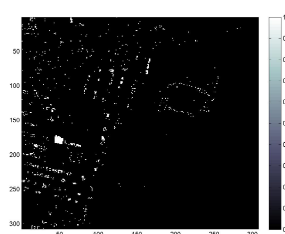

(j) TennisCrt

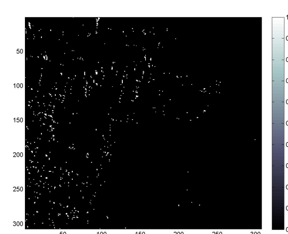

(k) PoolWater

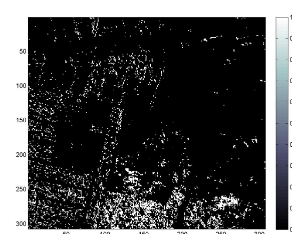

(l) ShadedVeg

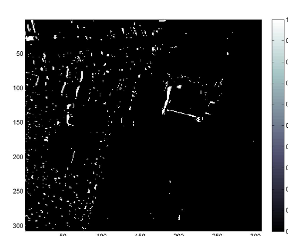

(m) ShadedPav

Figure 4.9: Urban trial 1 A individual class maps 10–22

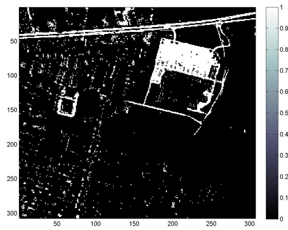

(a) AsphaltDrk

(b) AsphaltLgt

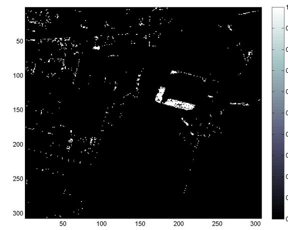

(c) Concrete01

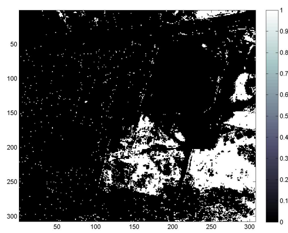

(d) VegPasture

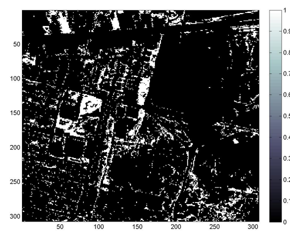

(e) VegGrass

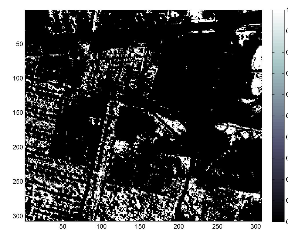

(f) VegTrees01

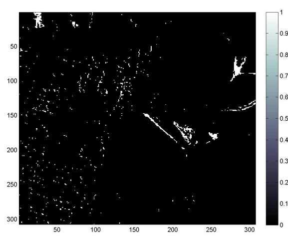

(g) Soil01

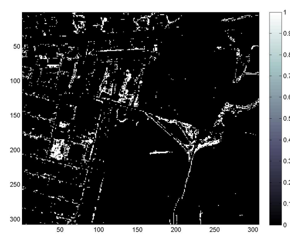

(h) Soil02

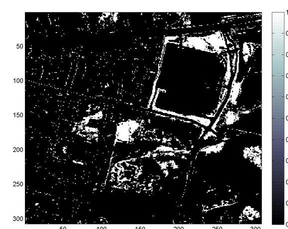

(i) Soil03Drk

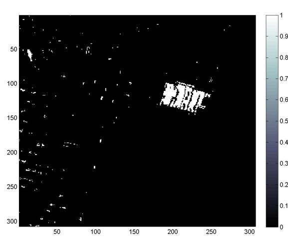

(j) Roof01Wal

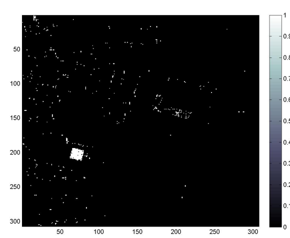

(k) Roof02A

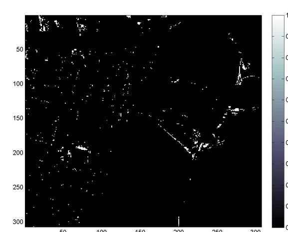

(l) Roof02BGvl

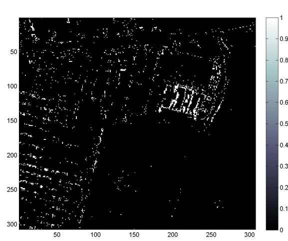

(m) Roof03LgtGray

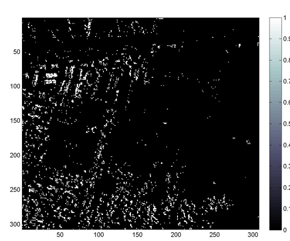

(n) Roof04DrkBrn

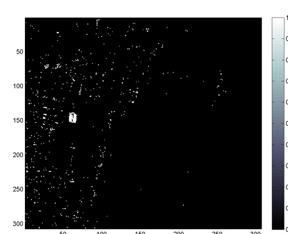

(o) Roof05AChurch

Figure 4.10: Urban trial 1 B individual class maps 1–15

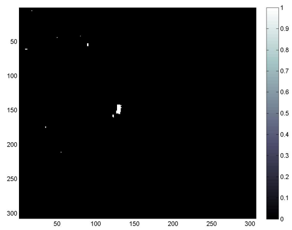

(a) Roof06School

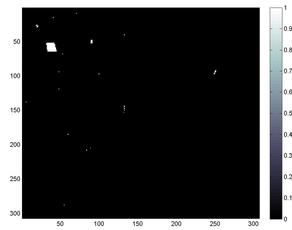

(b) Roof07Bright

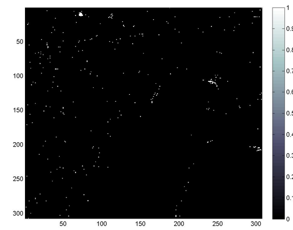

(c) Roof08BlueGrn

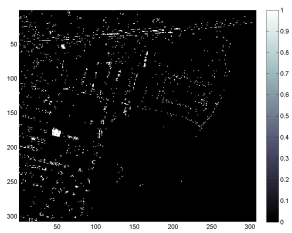

(d) TennisCrt

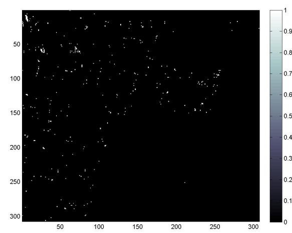

(e) PoolWater

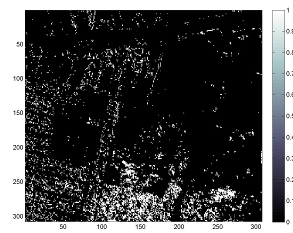

(f) ShadedVeg

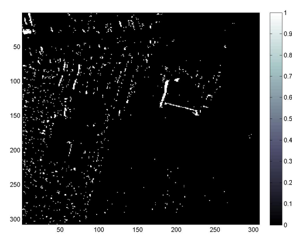

(g) ShadedPav

Figure 4.11: Urban trial 1 B individual class maps 16–22

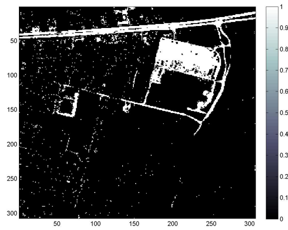

(a) AsphaltDrk

(b) AsphaltLgt

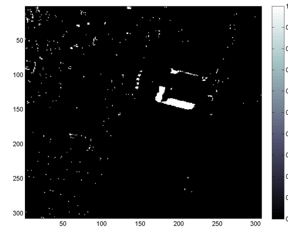

(c) Concrete01

(d) VegPasture

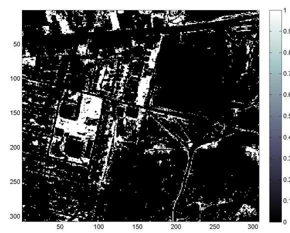

(e) VegGrass

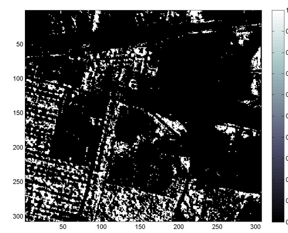

(f) VegTrees01

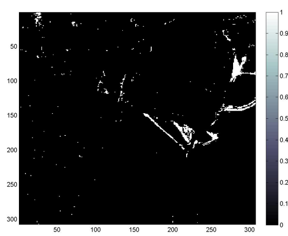

(g) Soil01

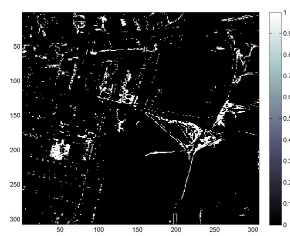

(h) Soil02

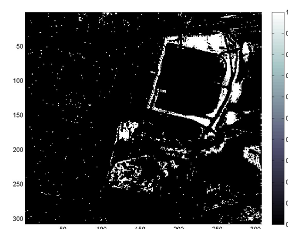

(i) Soil03Drk

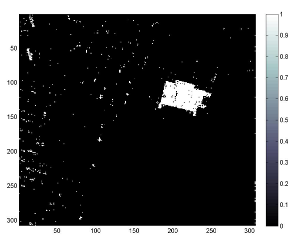

(j) Roof01Wal

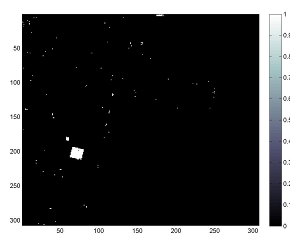

(k) Roof02A

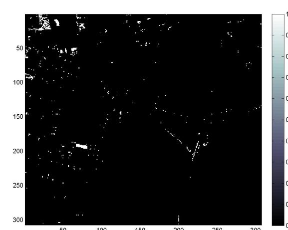

(l) Roof02BGvl

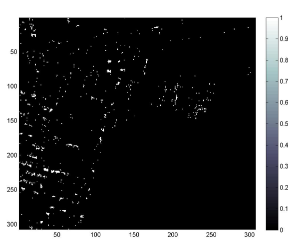

(m) Roof03LgtGray

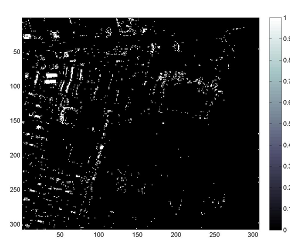

(n) Roof04DrkBrn

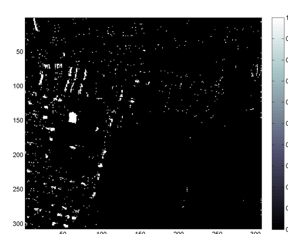

(o) Roof05AChurch

Figure 4.12: Urban trial 2 A individual class maps 1–15

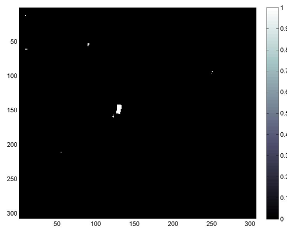

(a) Roof06School

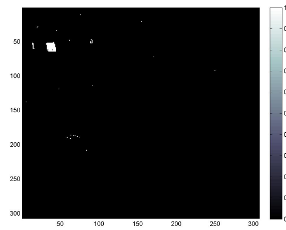

(b) Roof07Bright

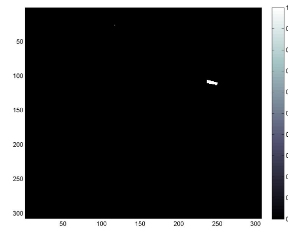

(c) Roof08BlueGrn

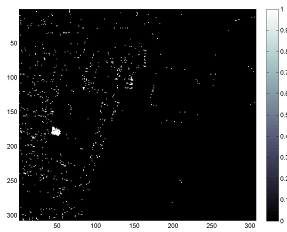

(d) TennisCrt

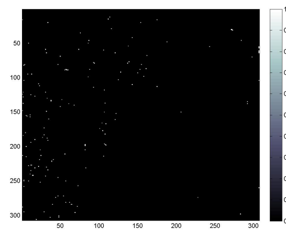

(e) PoolWater

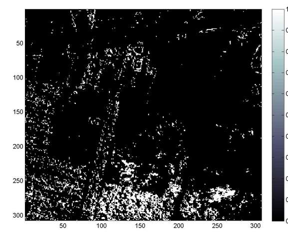

(f) ShadedVeg

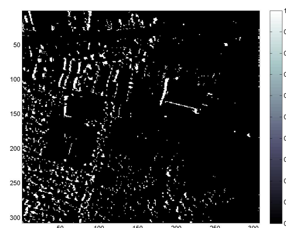

(g) ShadedPav

Figure 4.13: Urban trial 2 A individual class maps 16–22

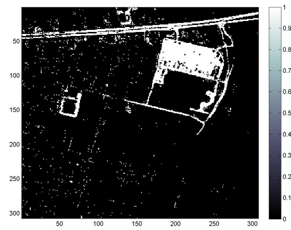

(a) AsphaltDrk

(b) AsphaltLgt

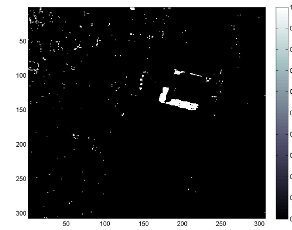

(c) Concrete01

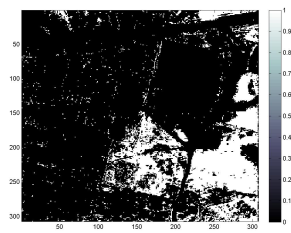

(d) VegPasture

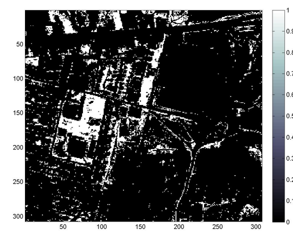

(e) VegGrass

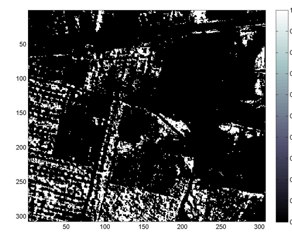

(f) VegTrees01

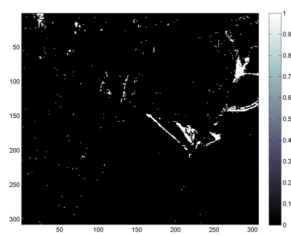

(g) Soil01

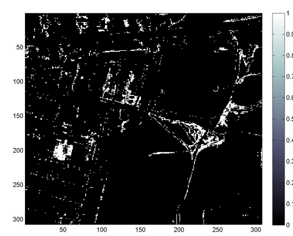

(h) Soil02

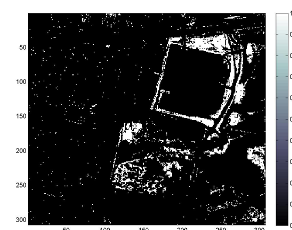

(i) Soil03Drk

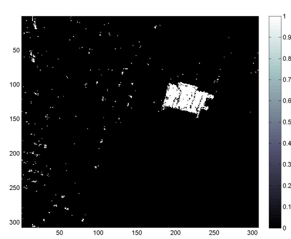

(j) Roof01Wal

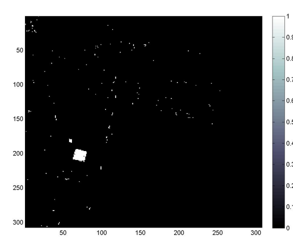

(k) Roof02A

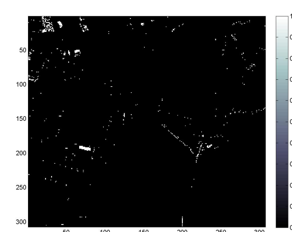

(l) Roof02BGvl

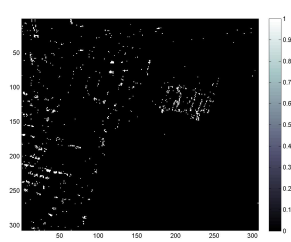

(m) Roof03LgtGray

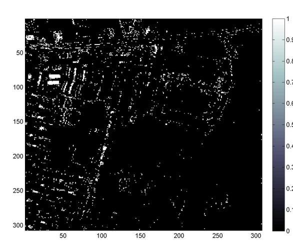

(n) Roof04DrkBrn

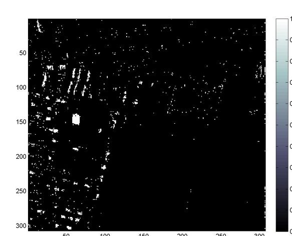

(o) Roof05AChurch

Figure 4.14: Urban trial 2 B individual class maps 1–15

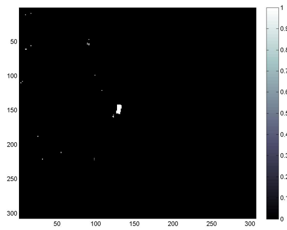

(a) Roof06School

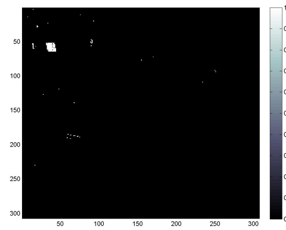

(b) Roof07Bright

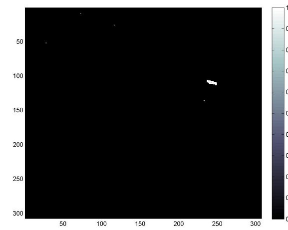

(c) Roof08BlueGrn

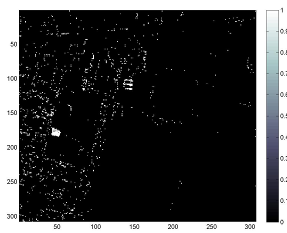

(d) TennisCrt

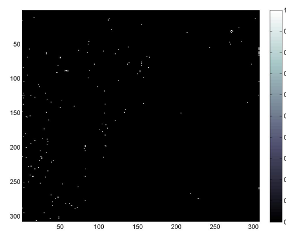

(e) PoolWater

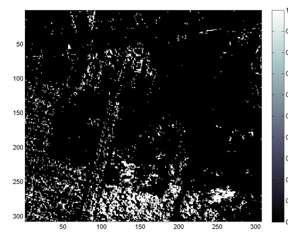

(f) ShadedVeg

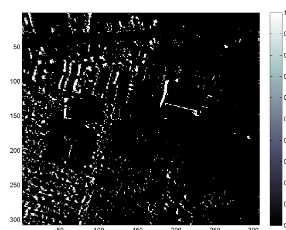

(g) ShadedPav

Figure 4.15: Urban trial 2 B individual class maps 16–22

4.1.3.7 Urban Coefficient Maps

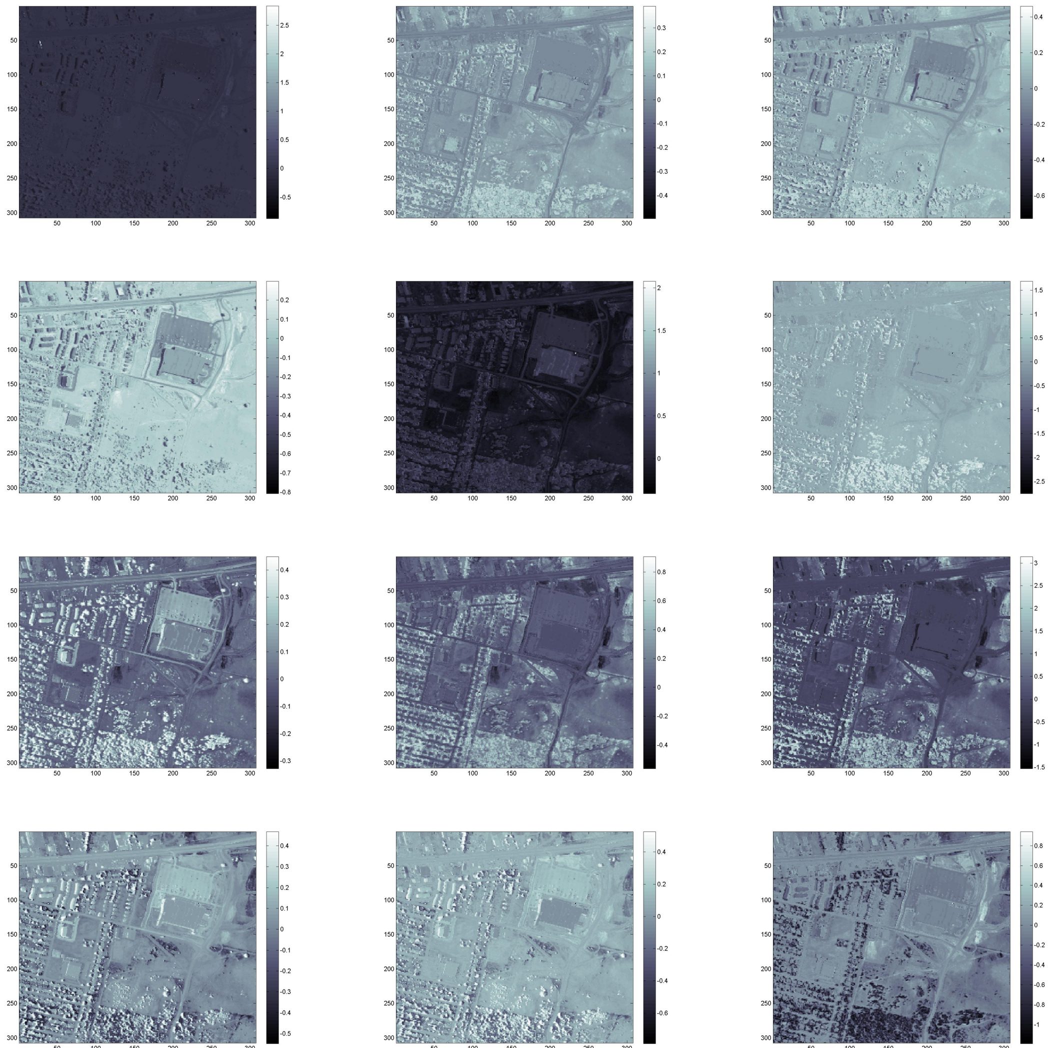

Figure 4.16: Urban trial 1 A canonical coefficients 1–12

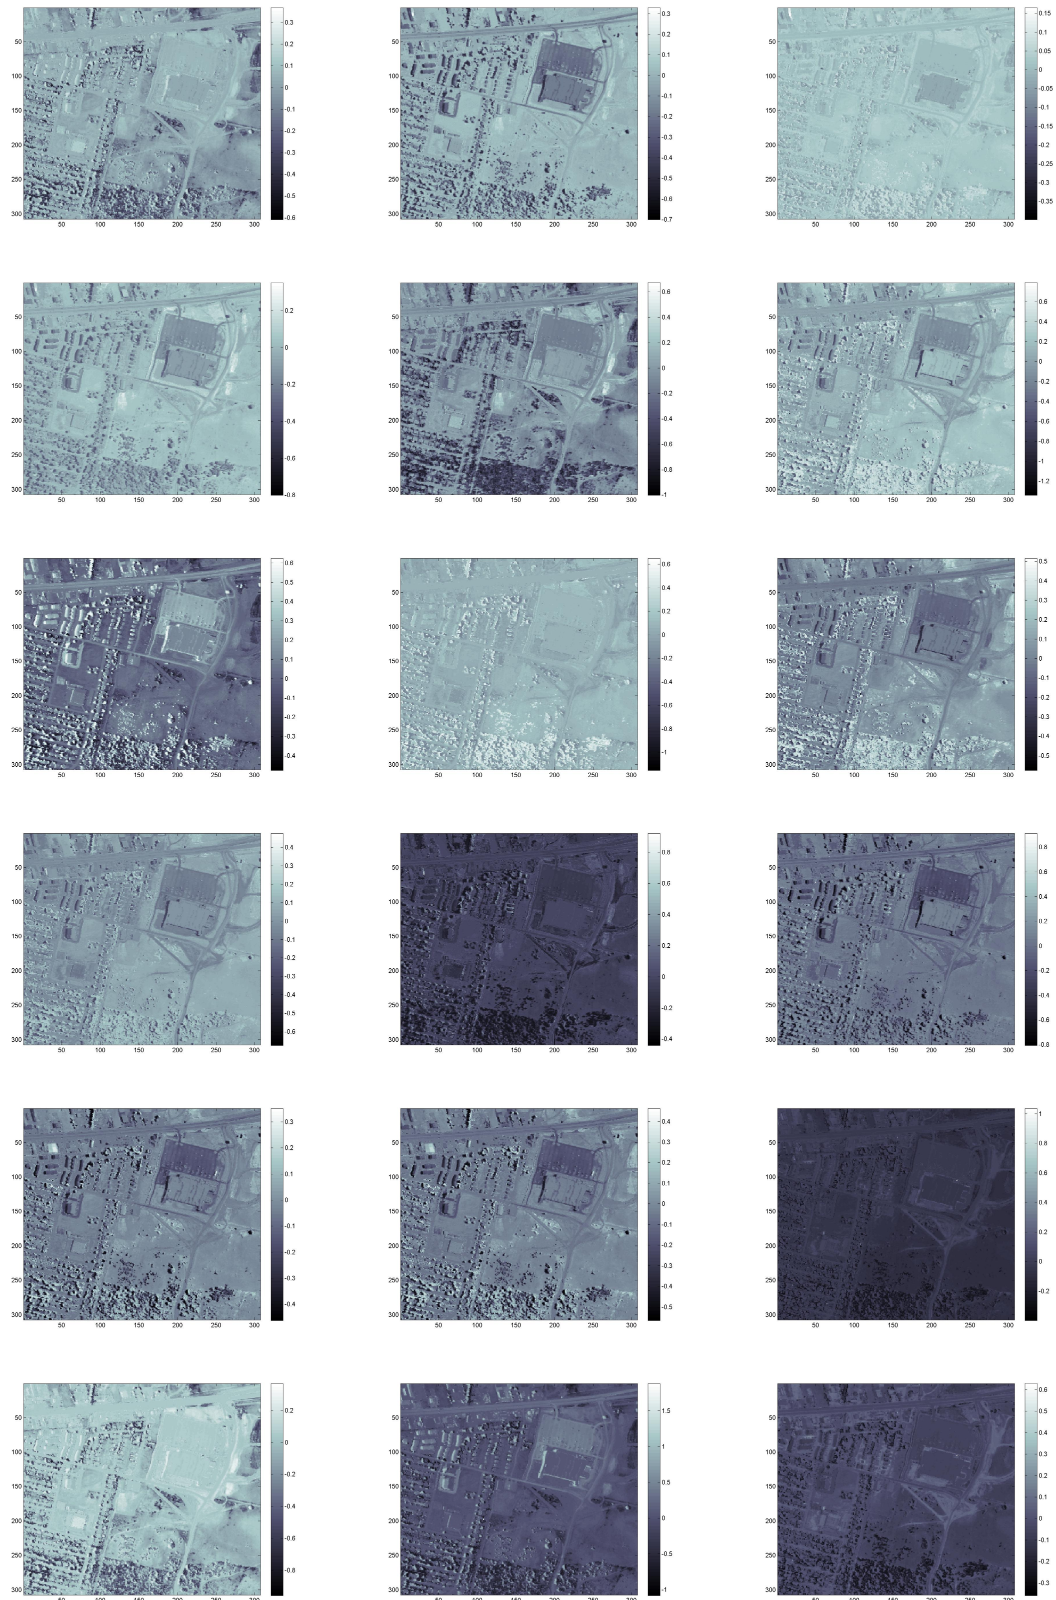

Figure 4.17: Urban trial 1 A canonical coefficients 13–30

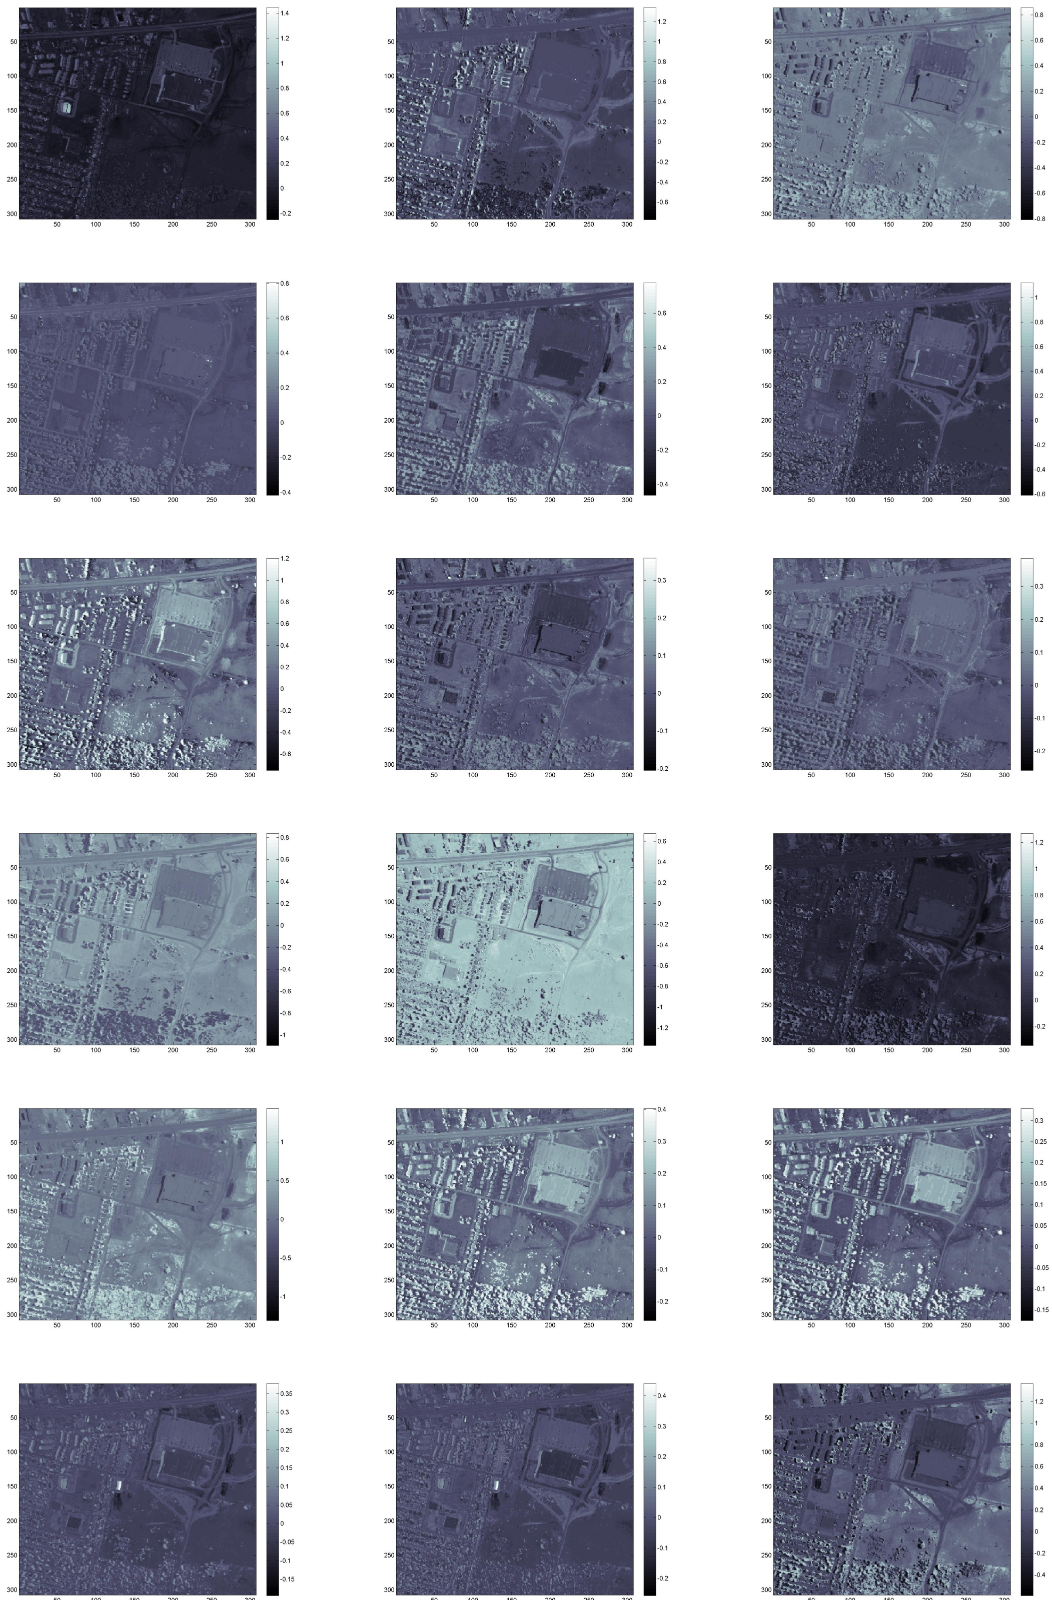

Figure 4.18: Urban trial 1 A canonical coefficients 31–48

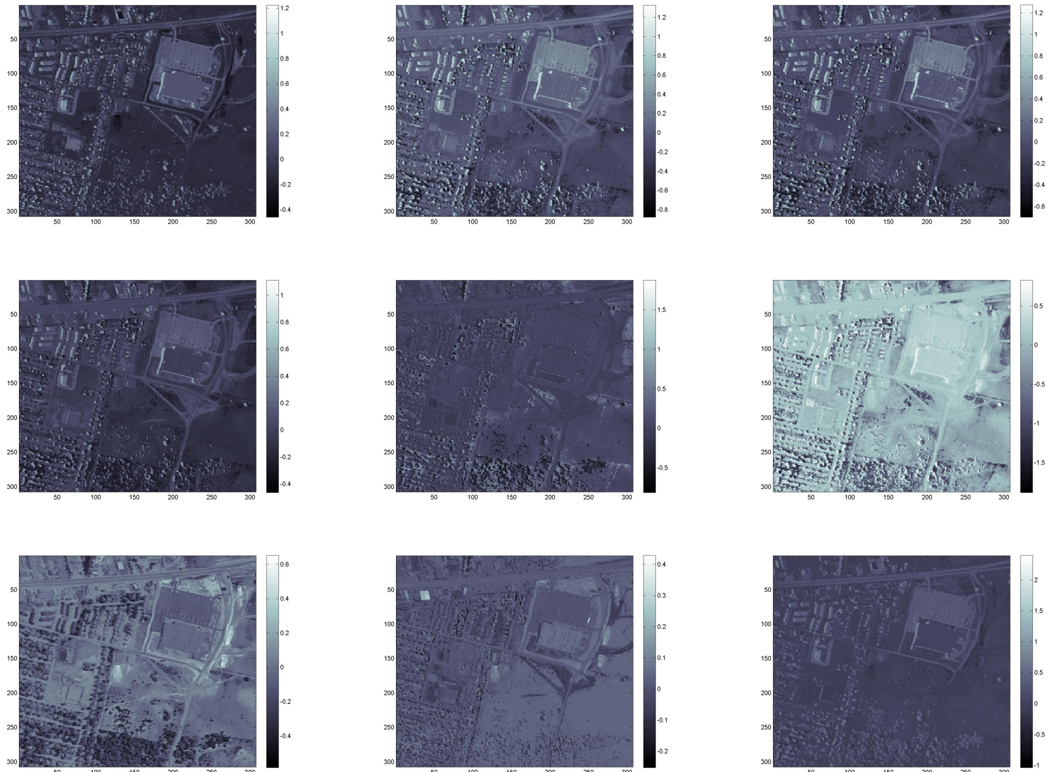

Figure 4.19: Urban trial 1 A canonical coefficients 49–57

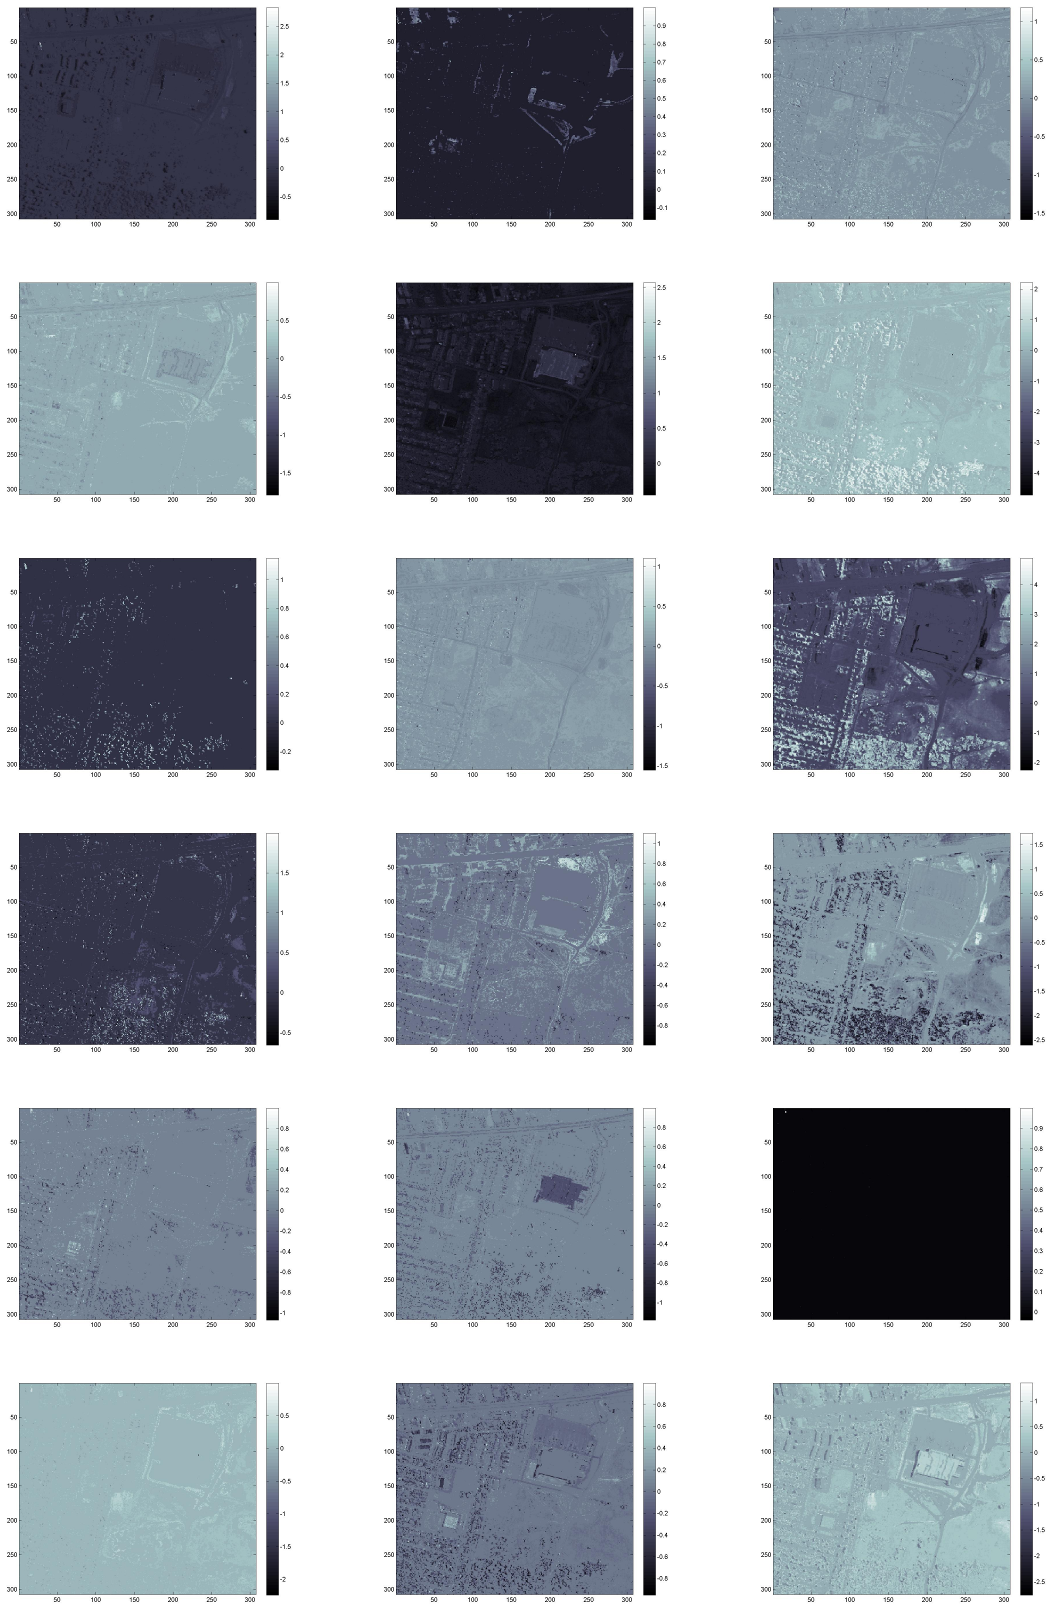

Figure 4.20: Urban trial 1 B sparse coefficients 1–18

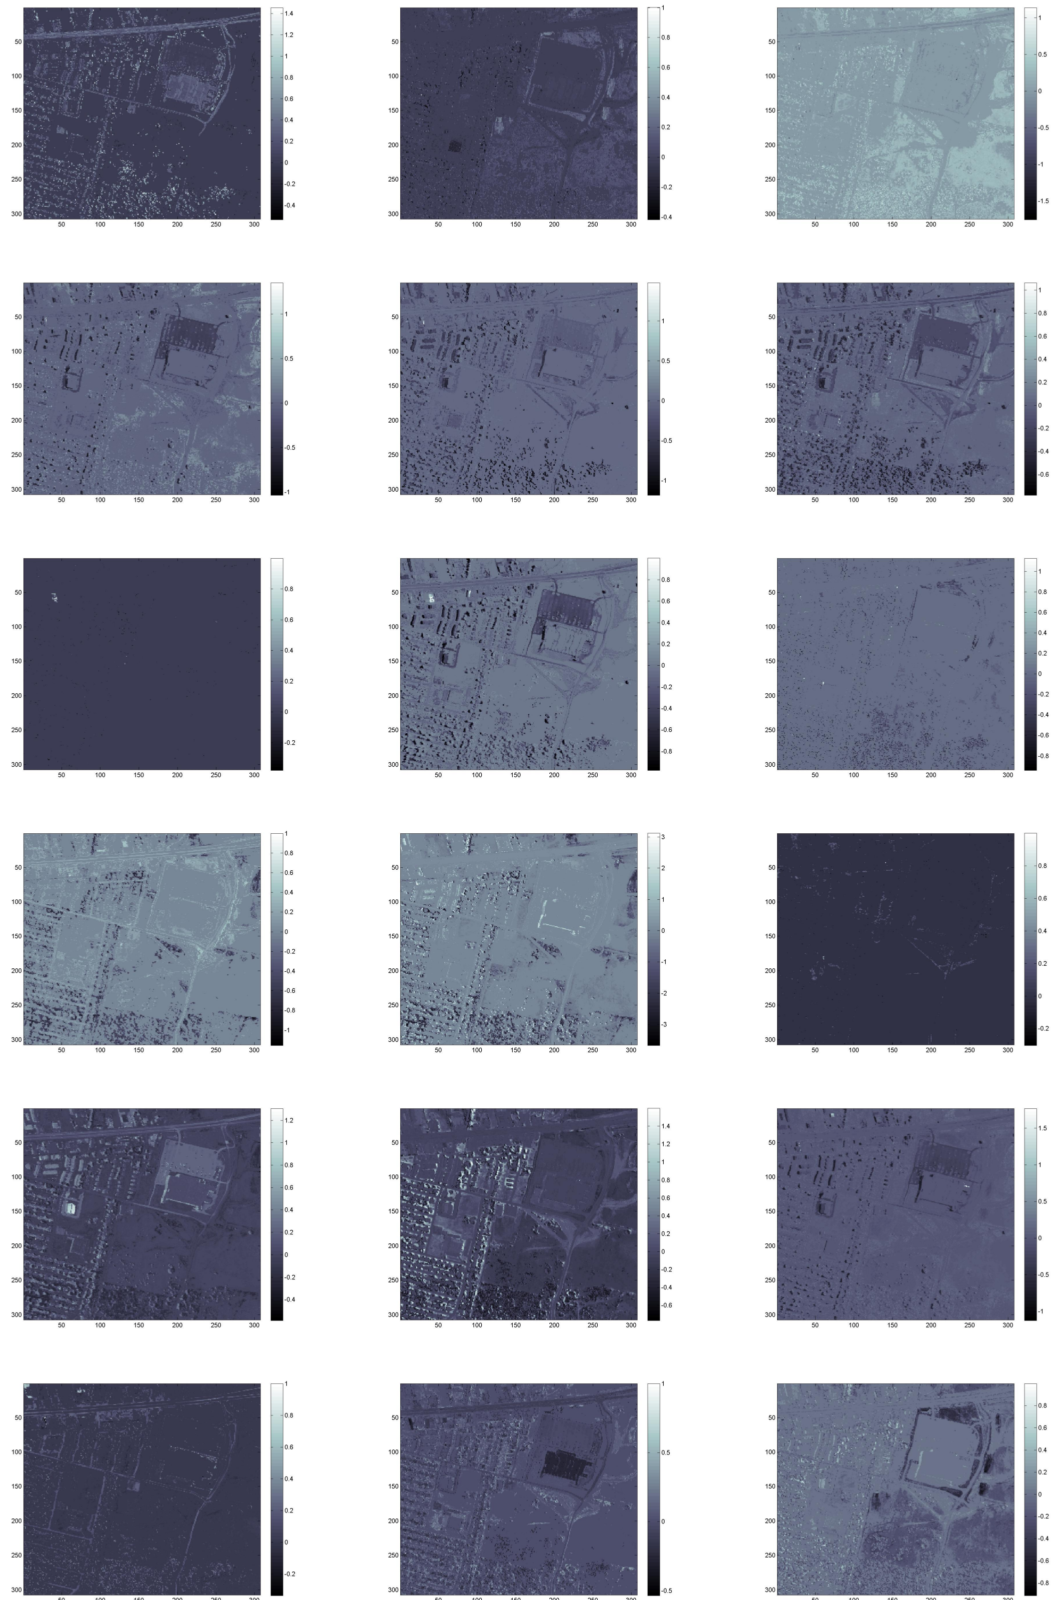

Figure 4.21: Urban trial 1 B sparse coefficients 19–36

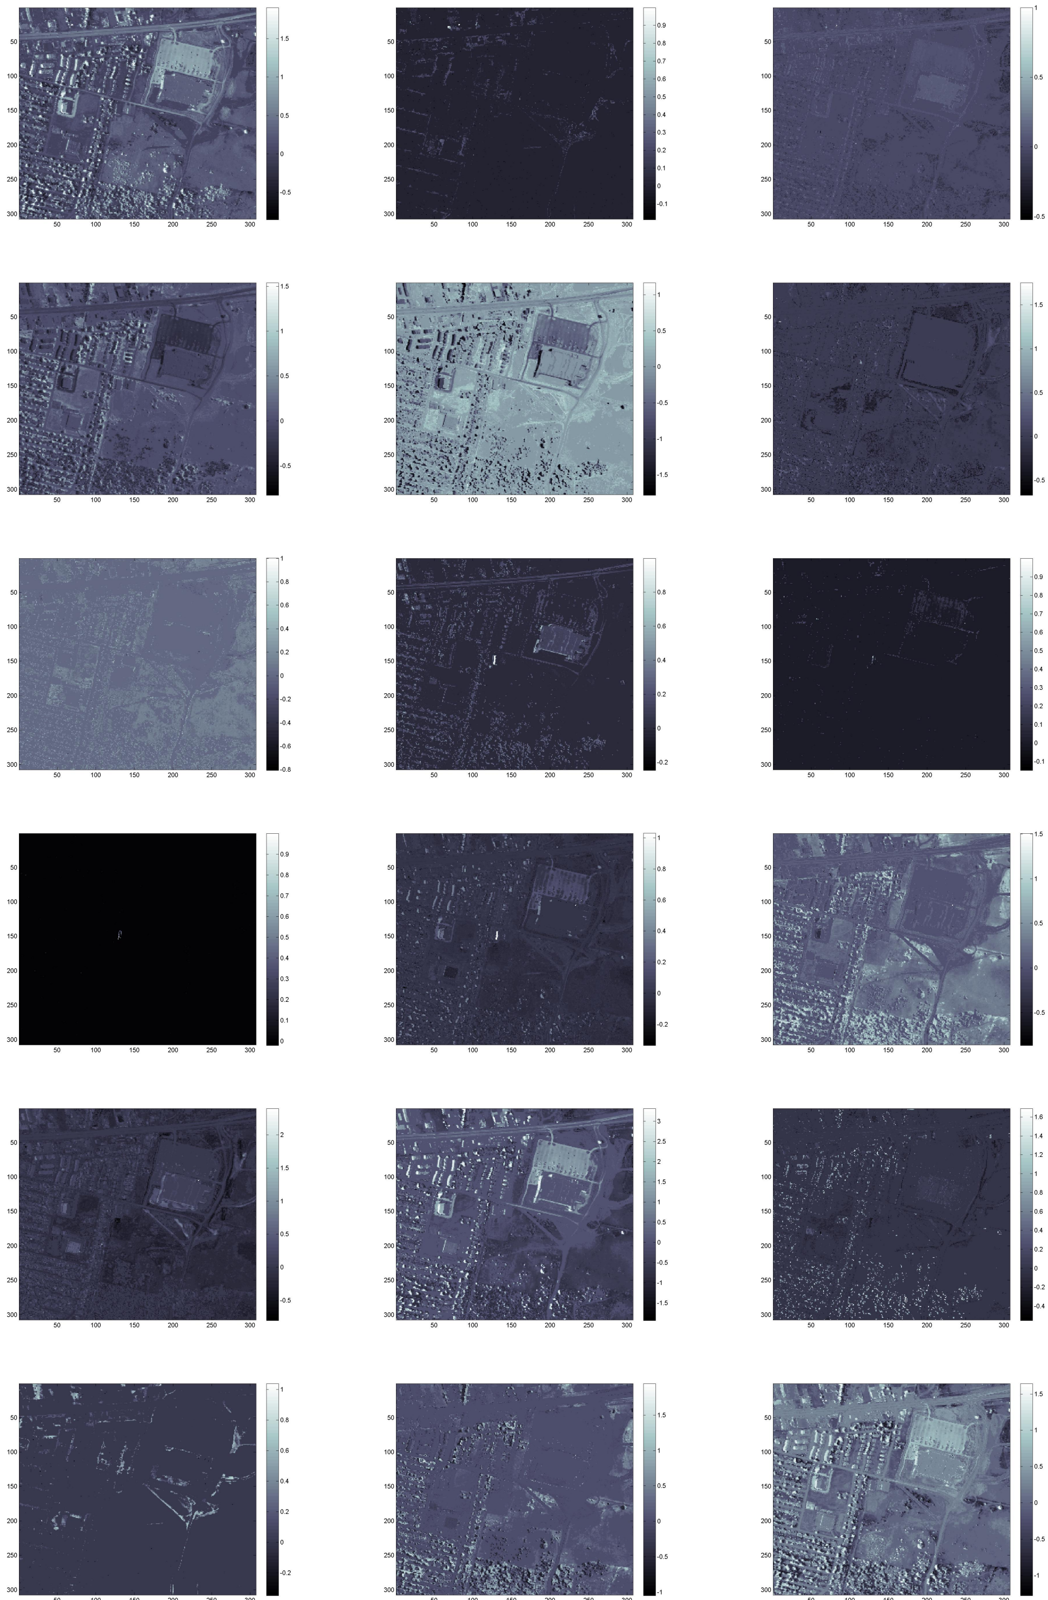

Figure 4.22: Urban trial 1 B sparse coefficients 37–54

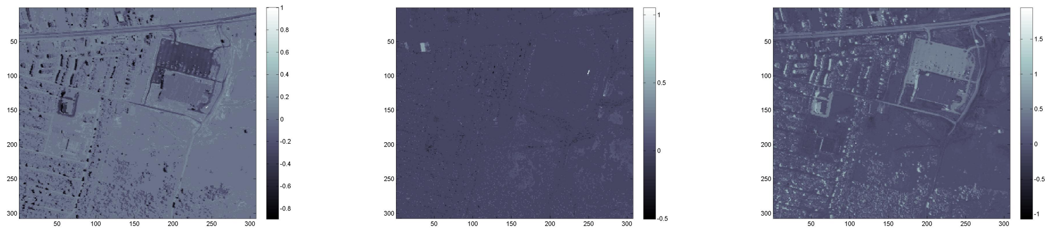

Figure 4.23: Urban trial 1 B sparse coefficients 55–57

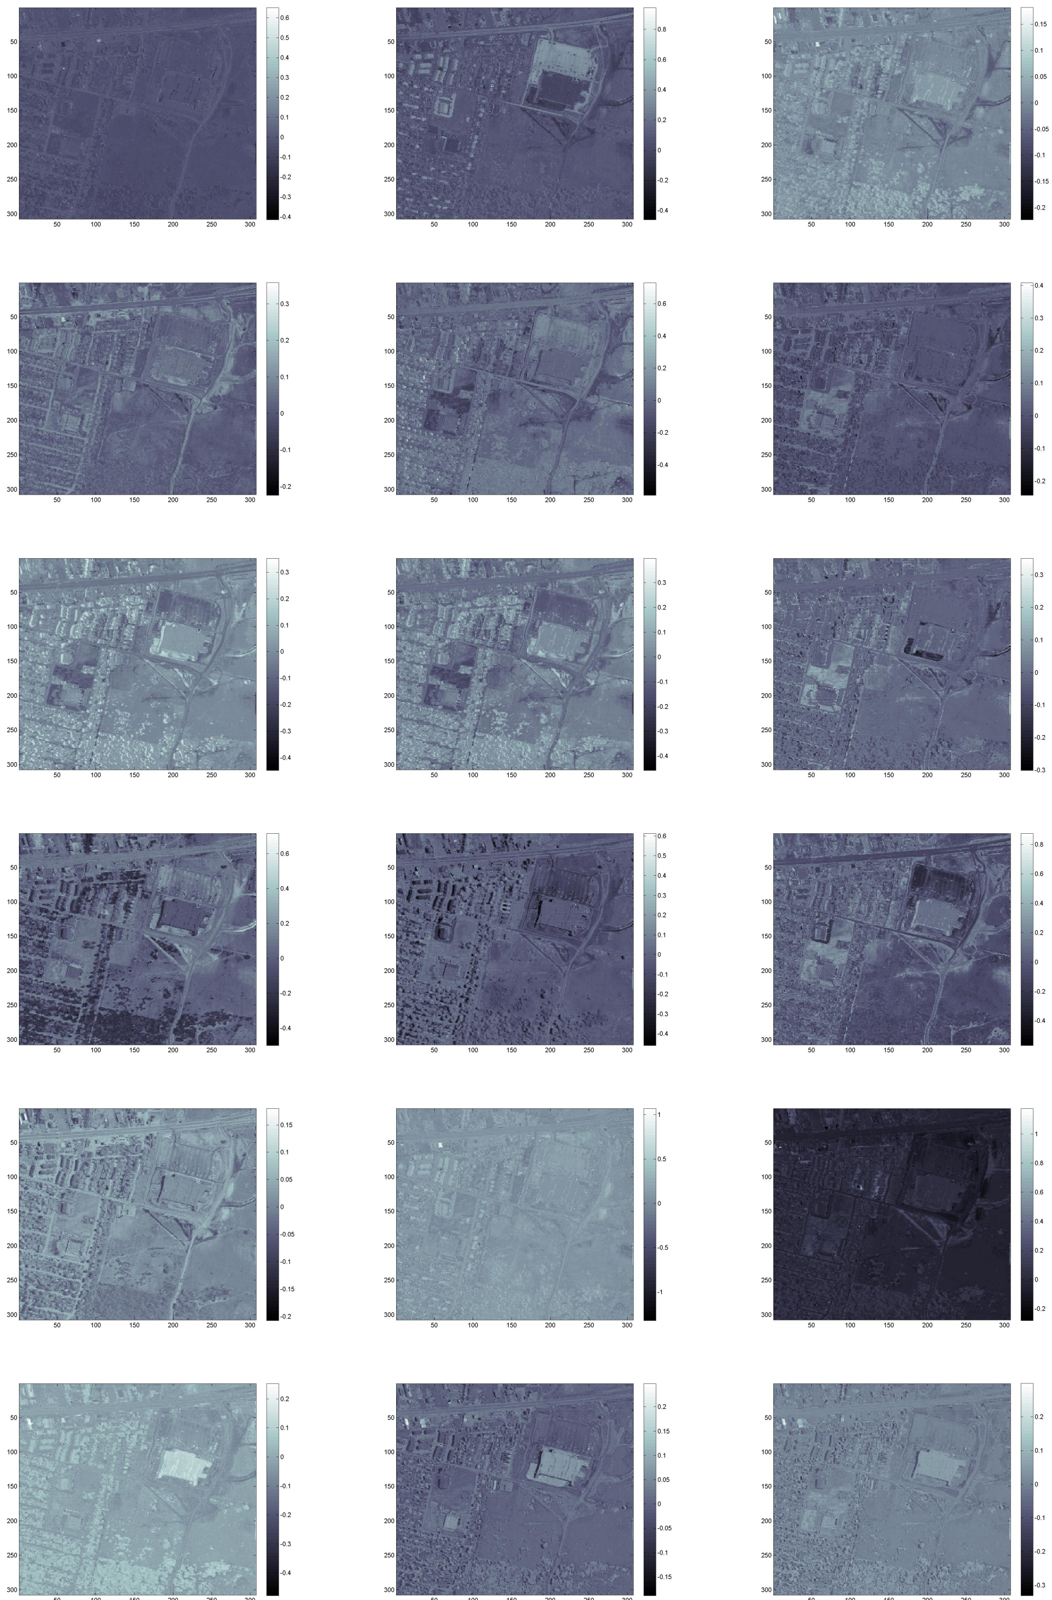

Figure 4.24: Urban trial 2 A canonical coefficients 1–18

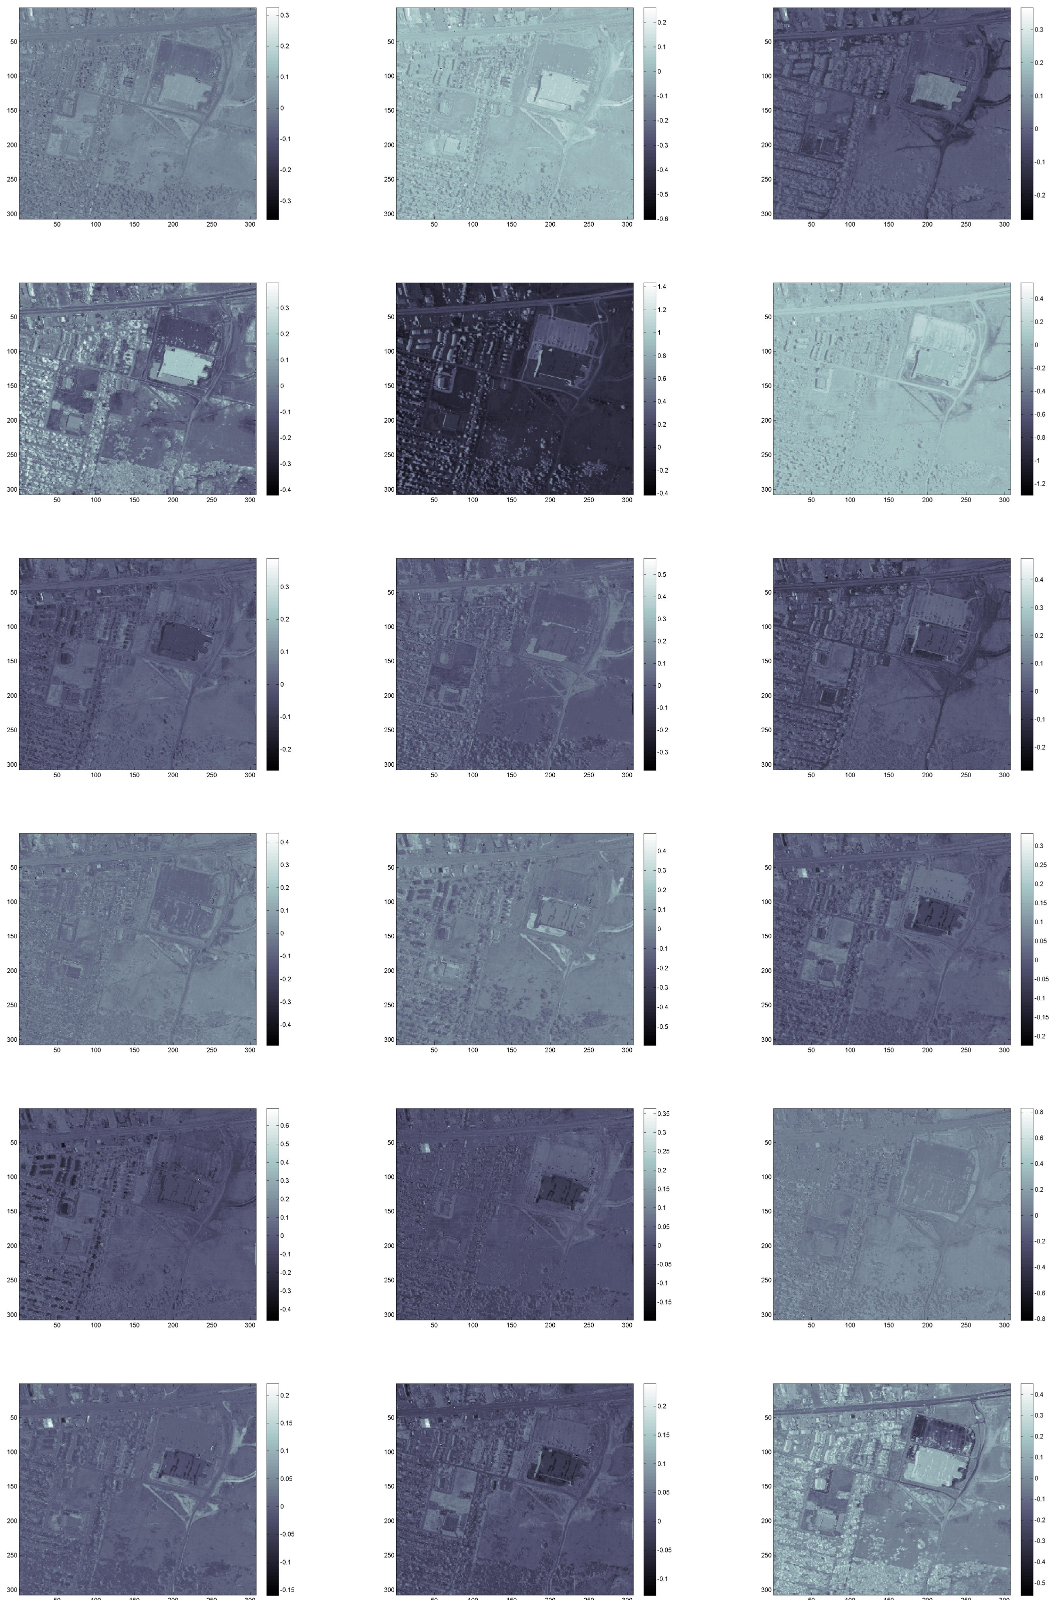

Figure 4.25: Urban trial 2 A canonical coefficients 19–36

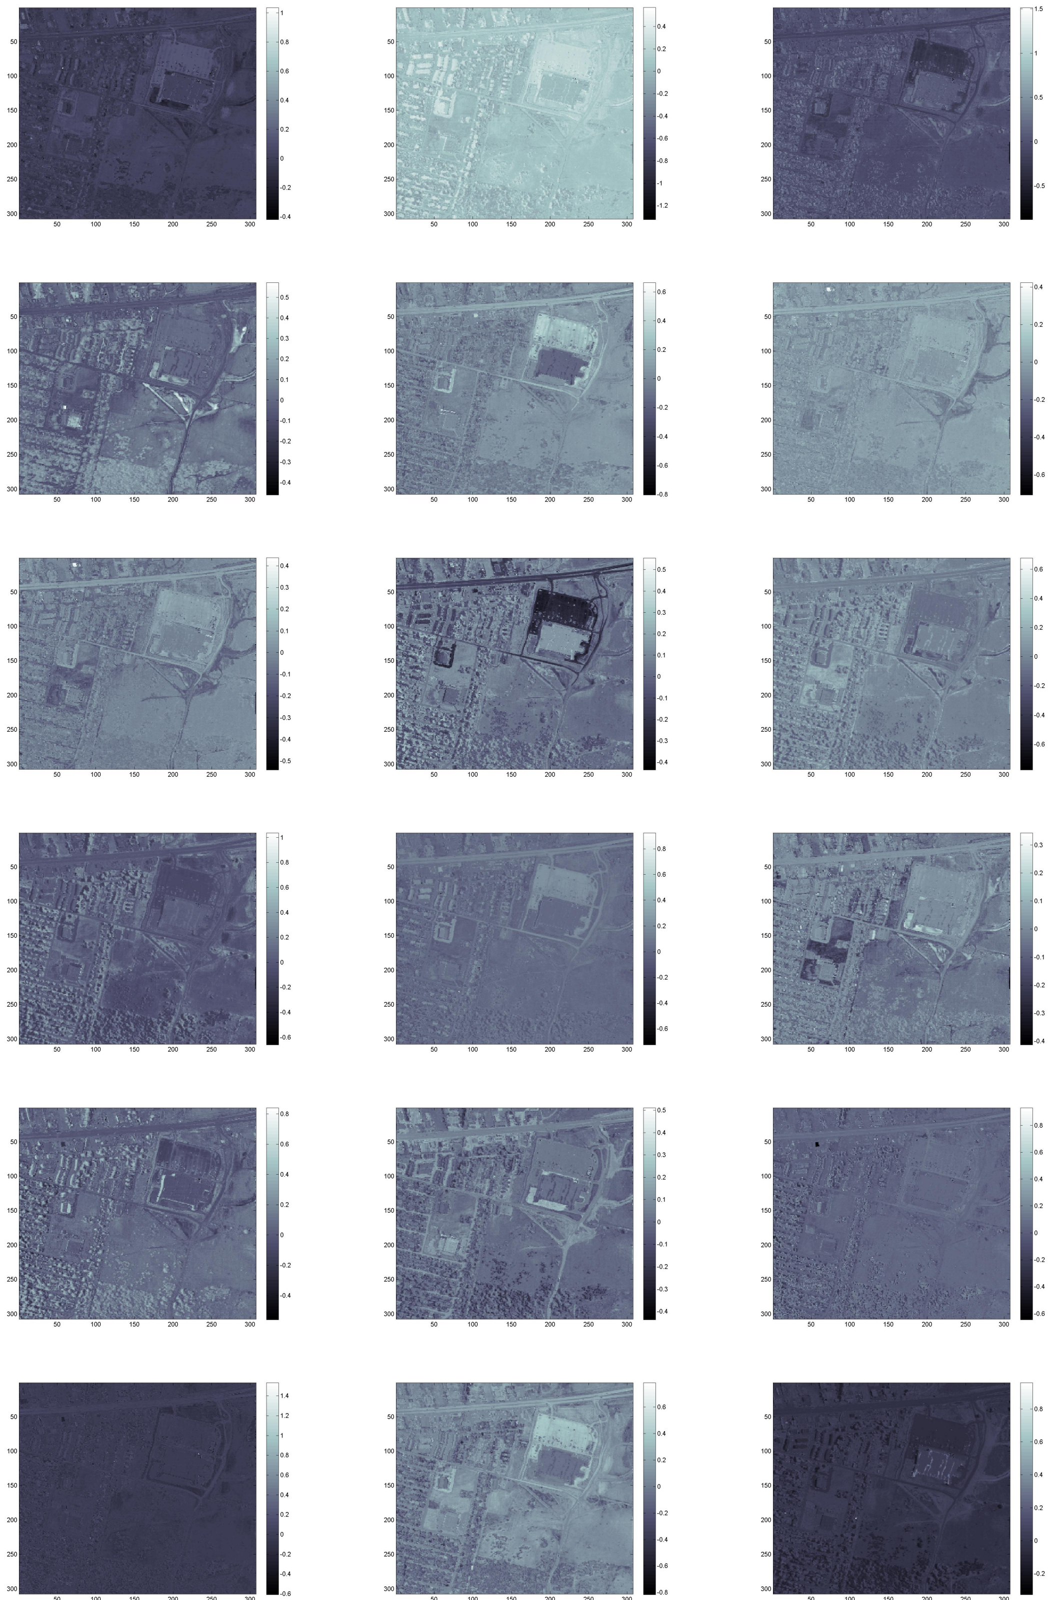

Figure 4.26: Urban trial 2 A canonical coefficients 37–54

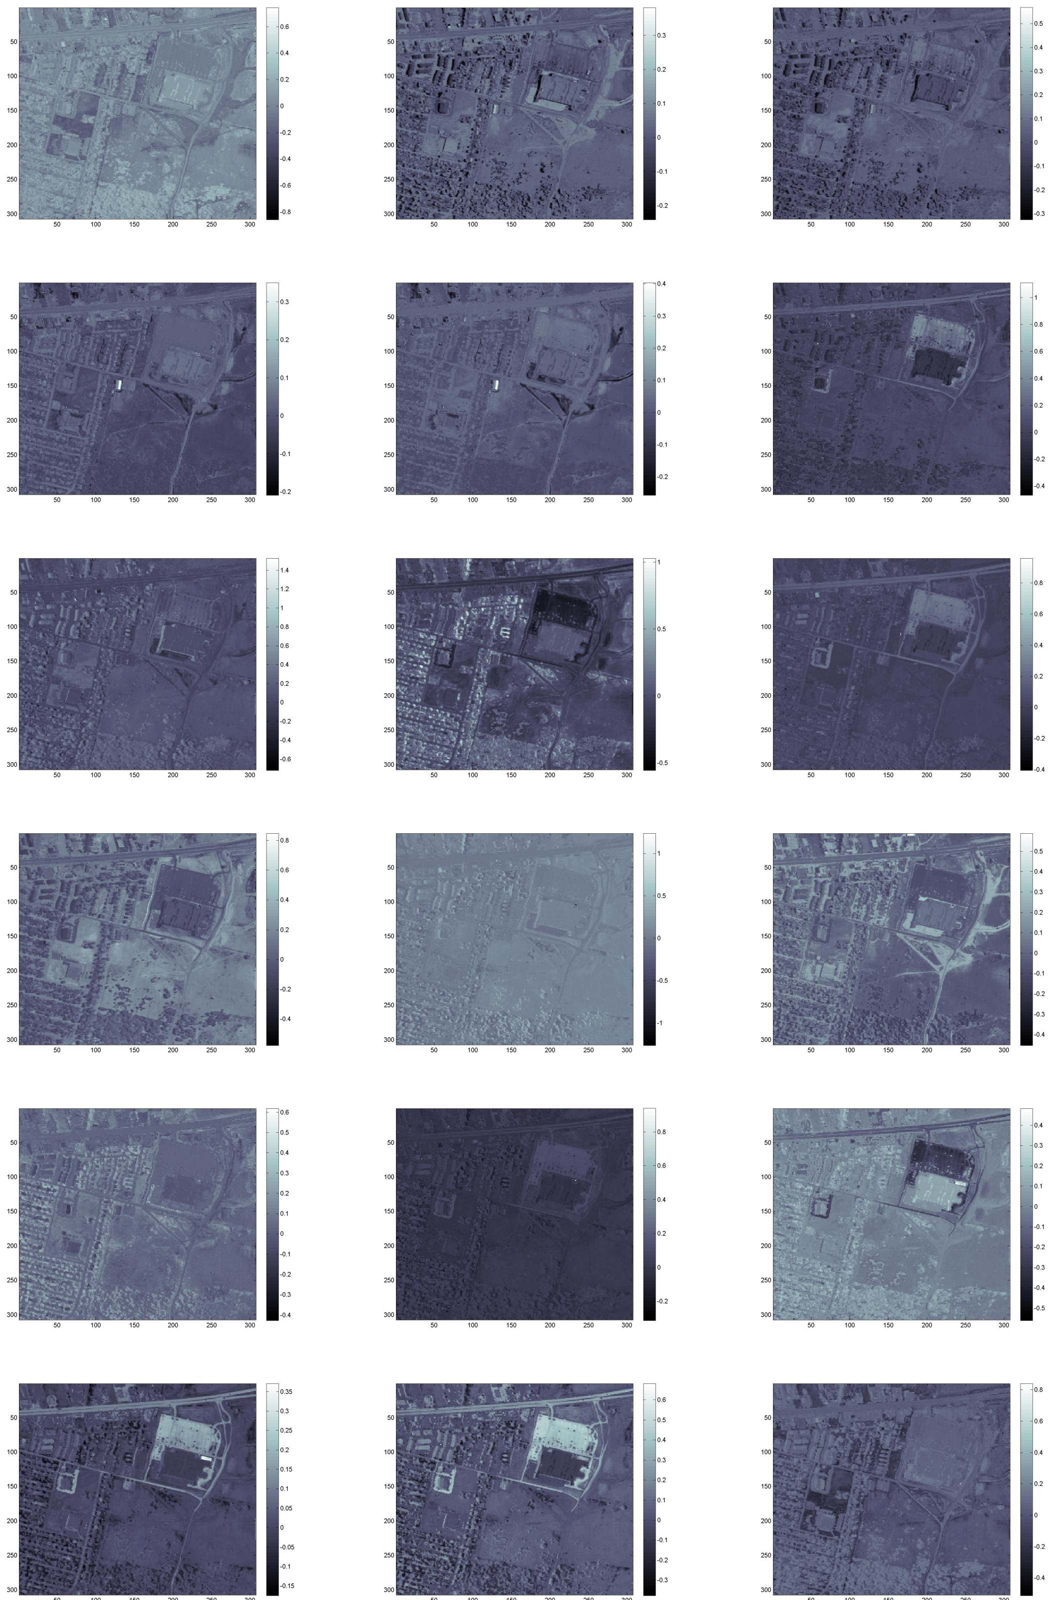

Figure 4.27: Urban trial 2 A canonical coefficients 55–72

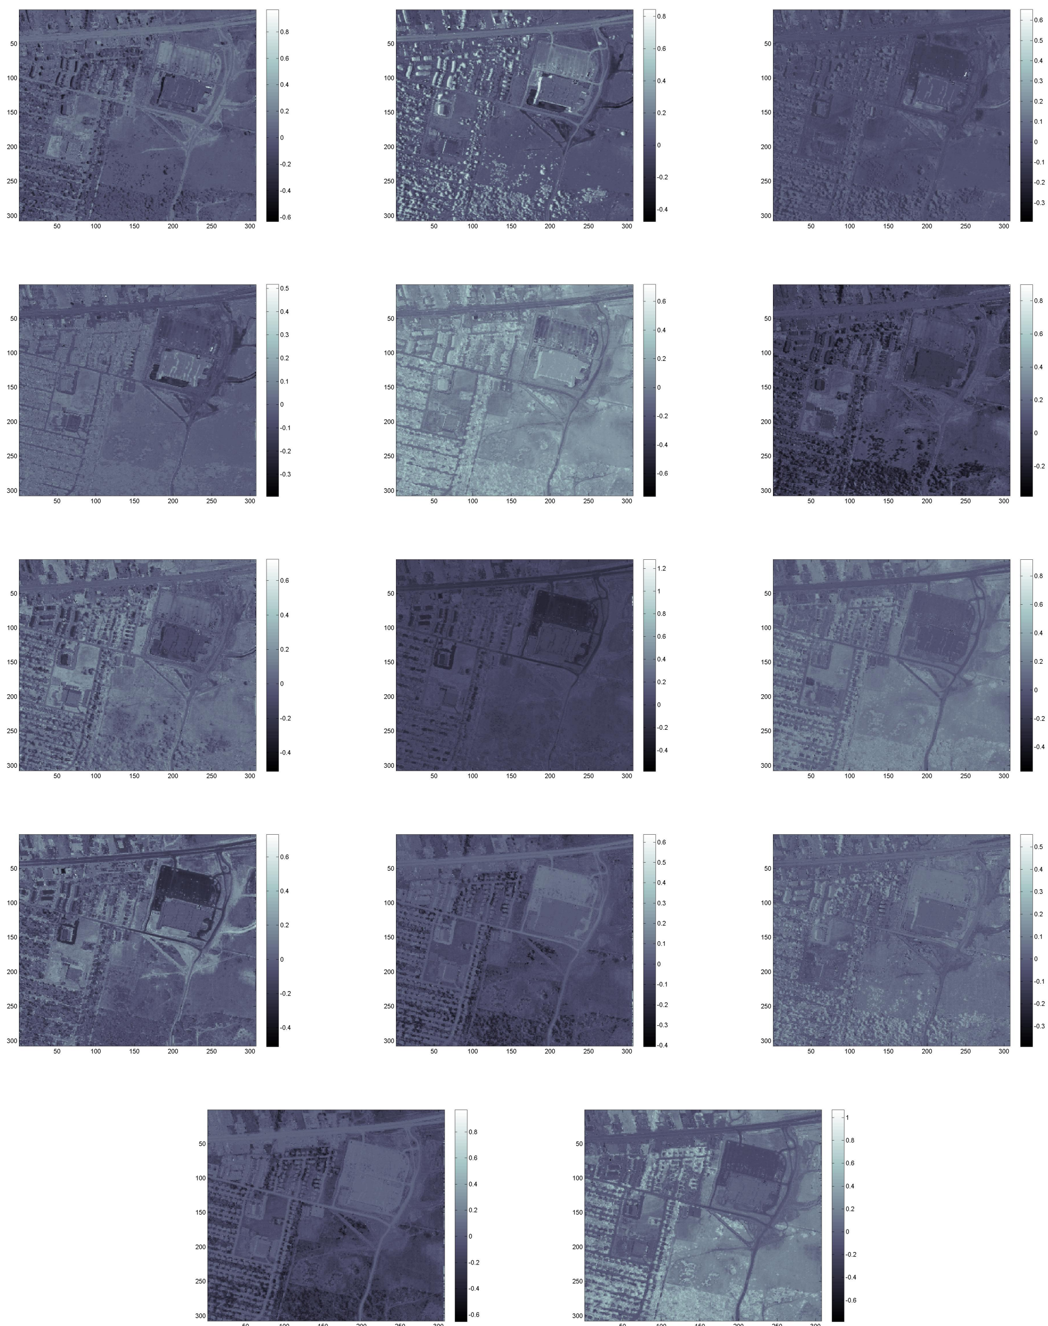

Figure 4.28: Urban trial 2 A canonical coefficients 73–86

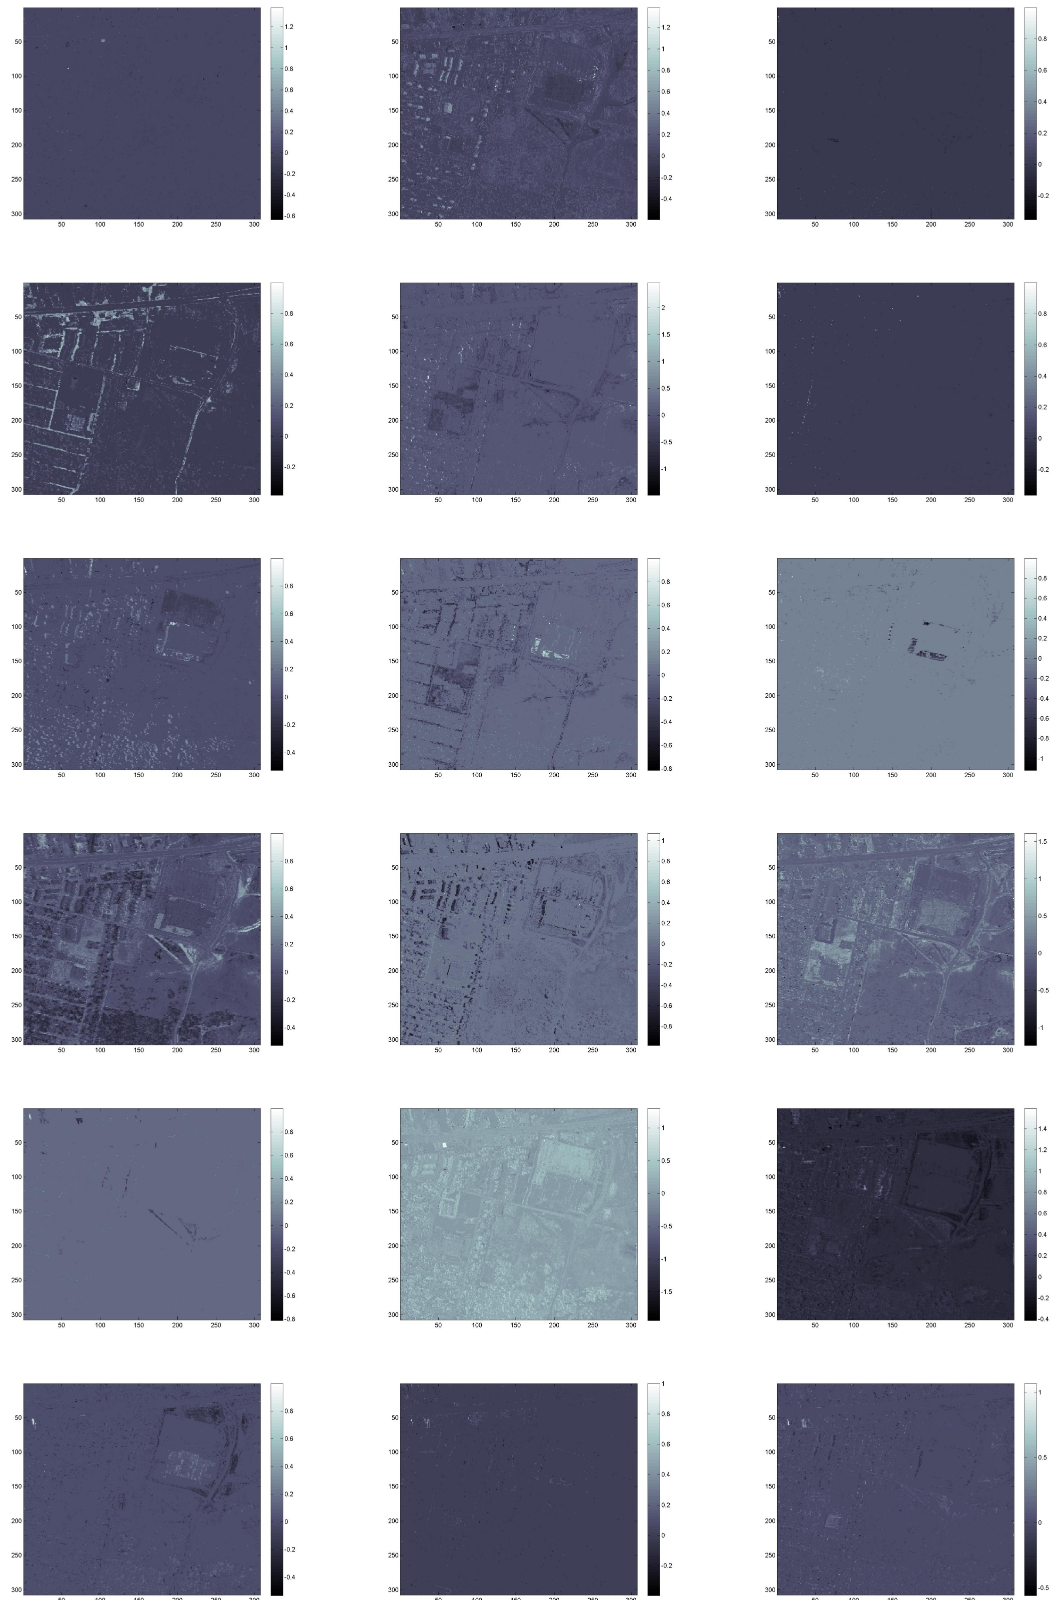

Figure 4.29: Urban trial 2 B sparse coefficients 1–18

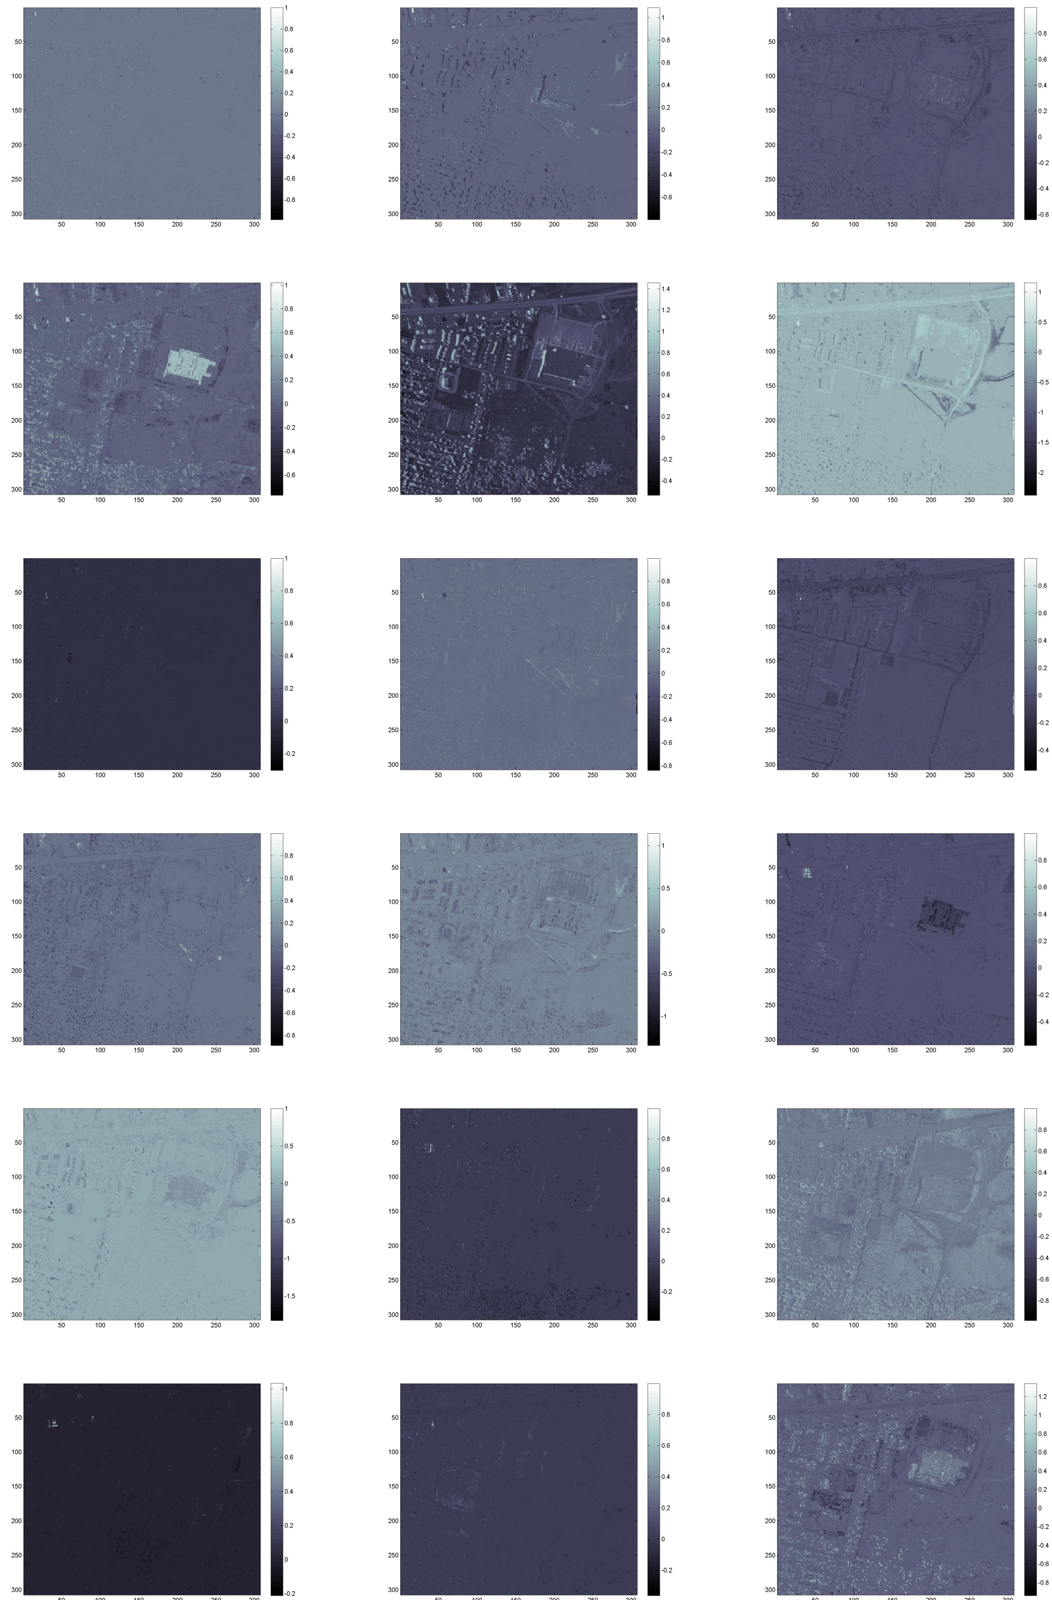

Figure 4.30: Urban trial 2 B sparse coefficients 19–36

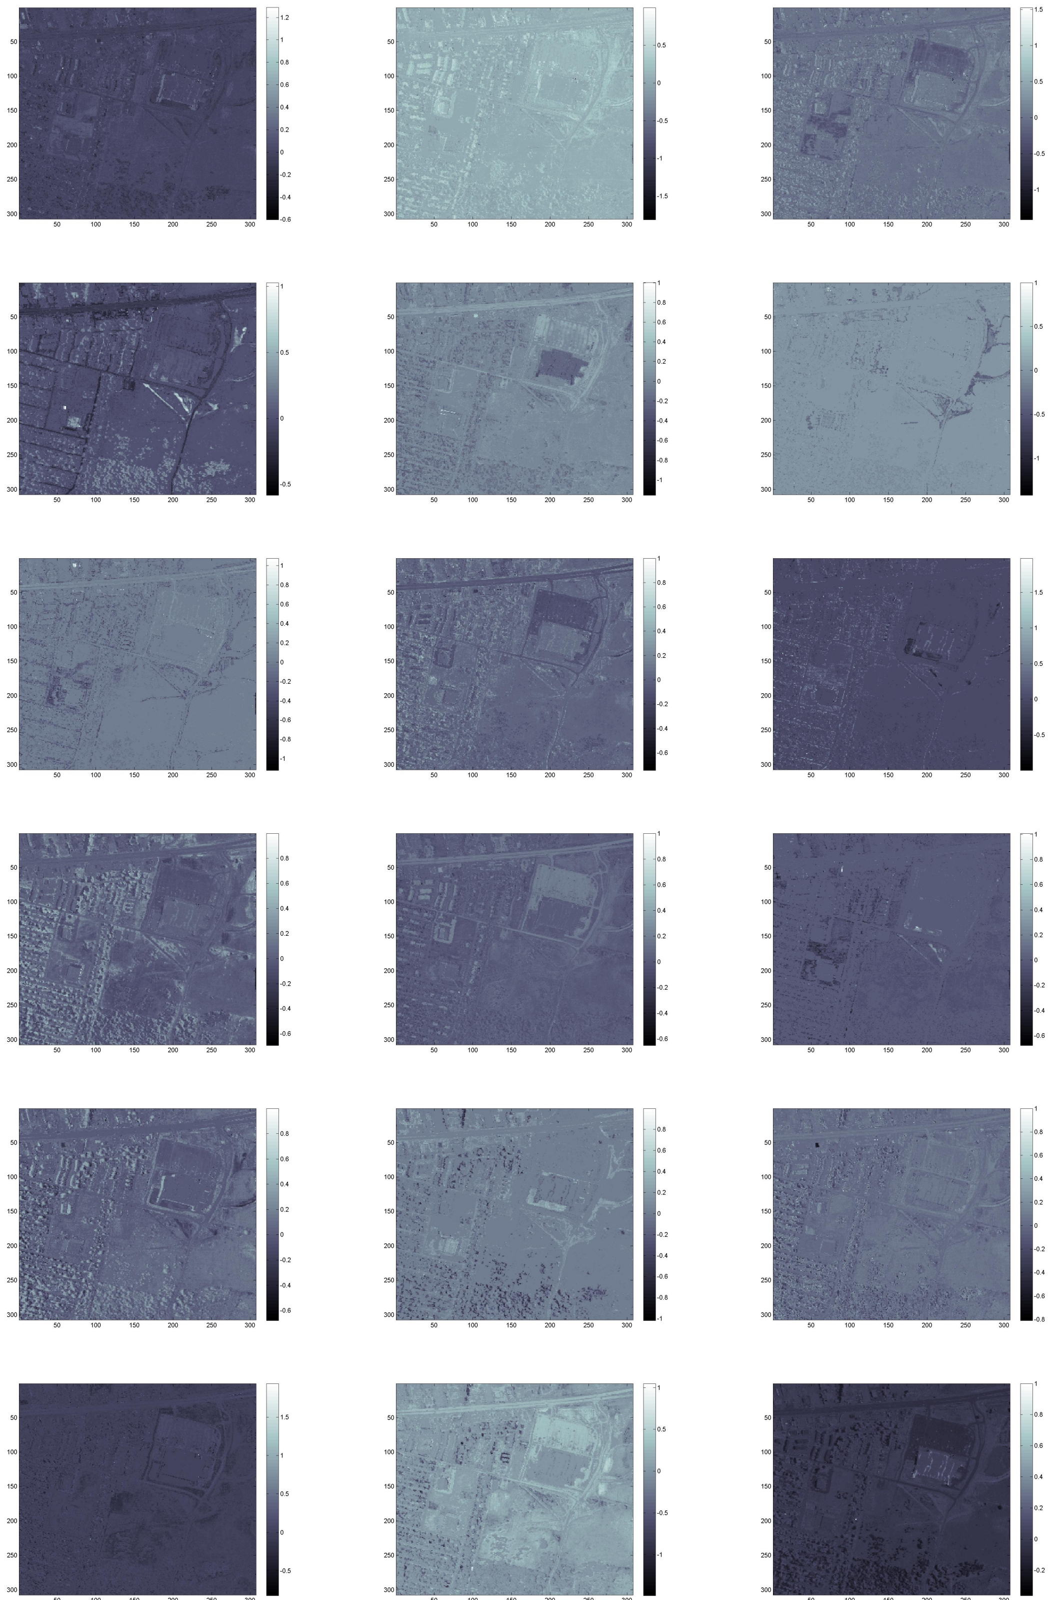

Figure 4.31: Urban trial 2 B sparse coefficients 37–54

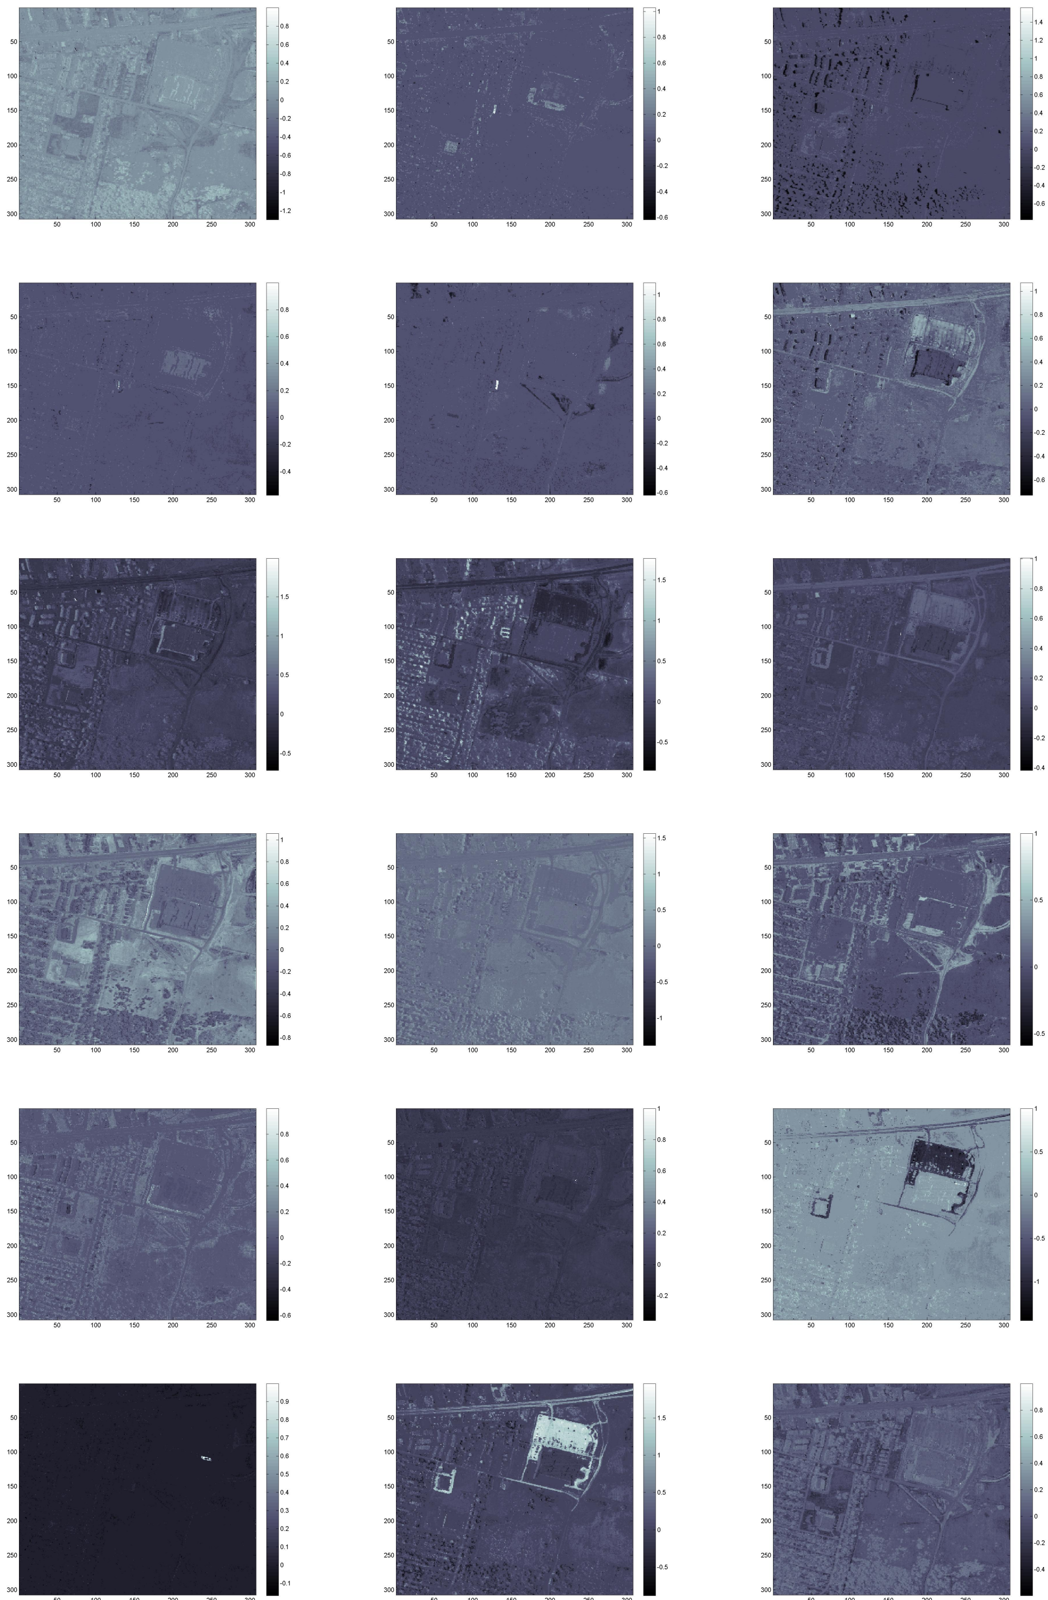

Figure 4.32: Urban trial 2 B sparse coefficients 55–72

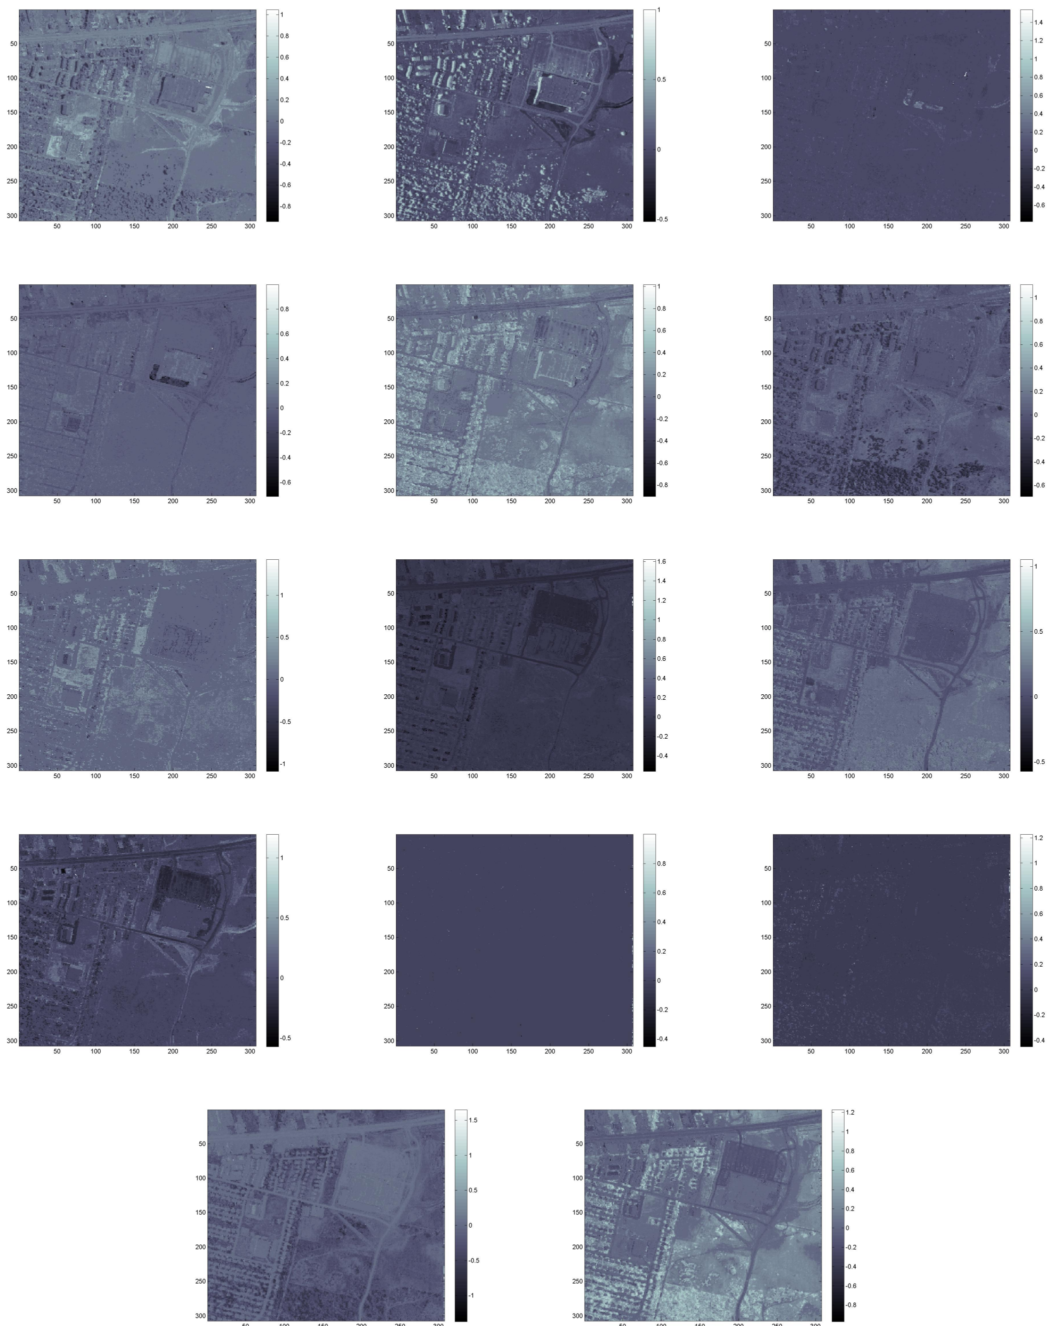

Figure 4.33: Urban trial 2 B sparse coefficients 73–86

4.1.4 Smith Island

4.1.4.1 Description of the Smith Island Data Set

The dimensions of the Smith Island data set are $679 \times 944 \times 110$: that is $679 * 944 = 640976$ pixels and 110 spectral bands. A pseudocolor image of the Smith Island data set is given in figure 4.34. There are 2743 ground truth pixels, broken into 22 distinct classes; these classes are:

1. phrag
2. scirpus
3. juncus
4. patens
5. distichlis
6. andropogon
7. ammophila
8. mud
9. alterniflora
10. borrichia
11. salicornia
12. iva

Figure 4.34: Pseudocolor image of Smith Island

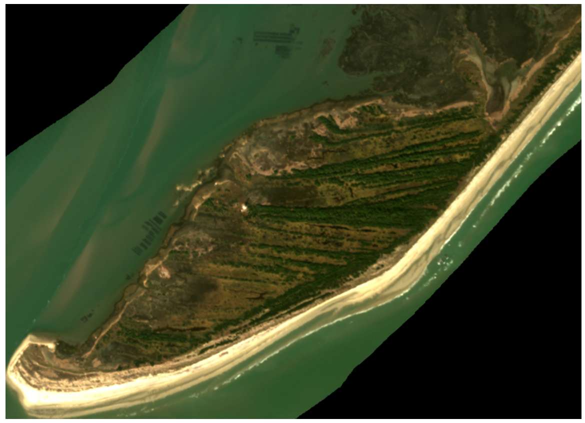

13. pine
14. hardwood
15. pond_water
16. sand
17. wrack
18. myrica
19. seaoats
20. typha
21. water_nshore
22. submerged_nets

4.1.4.2 Smith Island Trial 1

The results of Smith Island trial 1 were obtained with the following settings:

- Data set: X = Smith Island
- Kernel: LLE
- Number of neighbors: $k = 50$
- Number of samples: $\#X_{sam} = 40000$
- Frame construction: SVDD

The classification results for varying s and d and the canonical coefficients are displayed in figure 4.35. We highlight the following particular cases.

Smith Island Trial 1 A

- Number of reduced dimensions: $d = 21$
- Number of frame elements: $s = 69$
- Type of coefficients: canonical

Statistical results for Smith Island trial 1 A can be found in table 4.6. Figure 4.36 shows the class map for this trial, while figures 4.38 and 4.39 show the individual class maps. Figures 4.42, 4.43, 4.44, 4.45, and 4.46 show the coefficient maps for each of the frame elements.

Smith Island Trial 1 B

- Number of reduced dimensions: $d = 21$

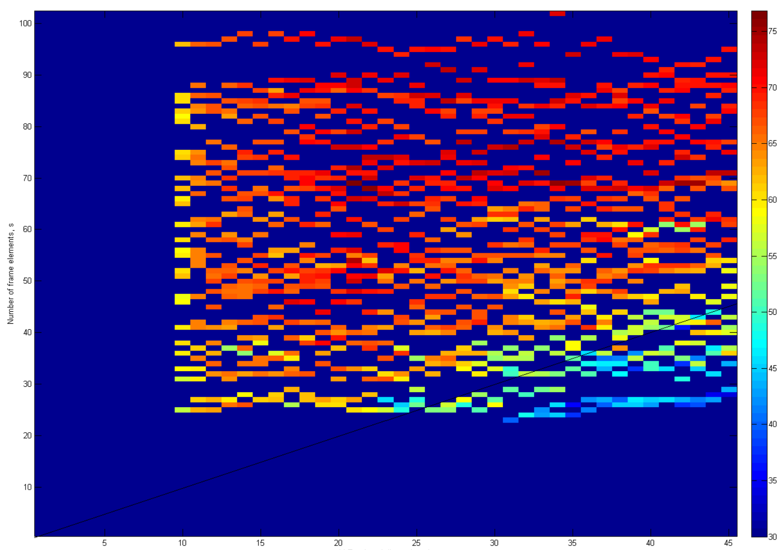

Figure 4.35: Smith Island trial 1 canonical coefficients classification results for varying s and d

Table 4.6: Smith trial 1 A ground truth results

	#	# correct	% correct	# false positives	# false negatives
phrag	196	138	70%	68	58
scirpus	246	155	63%	55	91
juncus	184	116	63%	33	68
patens	66	57	86%	33	9
distichlis	97	90	93%	18	7
andropogon	57	38	67%	9	19
ammophila	32	25	78%	29	7
mud	70	63	90%	25	7
alterniflora	200	182	91%	60	18
borrichia	90	84	93%	24	6
salicornia	76	58	76%	3	18
iva	58	49	84%	51	9
pine	166	134	81%	59	32
hardwood	328	193	59%	41	135
pond_water	105	69	66%	3	36
sand	159	157	99%	0	2
wrack	144	97	67%	11	47
myrica	167	132	79%	54	35
seaoats	18	13	72%	0	5
typha	44	18	41%	59	26
water_nshore	206	206	100%	0	0
submerged_nets	34	34	100%	0	0
Total	2743	2108	77%	635	635

- Number of frame elements: $s = 69$
- Type of coefficients: sparse

Statistical results for Smith Island trial 1 B can be found in table 4.7. Figure 4.37 shows the class map for this trial, while figures 4.40 and 4.41 show the individual class maps. Figures 4.47, 4.48, 4.49, and 4.50 show the coefficient maps for each of the frame elements.

Table 4.7: Smith trial 1 B ground truth results

	#	# correct	% correct	# false positives	# false negatives
phrag	196	129	67%	64	67
scirpus	246	160	65%	45	86
juncus	184	112	61%	54	72
patens	66	54	82%	31	12
distichlis	97	86	87%	35	11
andropogon	57	39	68%	6	18
ammophila	32	25	78%	36	7
mud	70	64	91%	23	6
alterniflora	200	190	95%	69	10
borrichia	90	85	94%	6	5
salicornia	76	57	75%	2	19
iva	58	32	55%	52	26
pine	166	112	67%	75	54
hardwood	328	189	58%	53	139
pond_water	105	70	67%	1	35
sand	159	159	100%	0	0
wrack	144	95	66%	9	49
myrica	167	107	64%	70	60
seaoats	18	13	72%	0	5
typha	44	19	43%	75	25
water_nshore	206	206	100%	0	0
submerged_nets	34	34	100%	0	0
Total	2743	2037	74%	706	706

4.1.4.3 Smith Island Competing Results

Table 4.8 contains the overall results of the competing Smith Island results. We note that the LLE and SVDD results were obtained at the following points:

- LLE only (trial 1): $d = 43$
- SVDD (both coefficient cubes): $s = 8$

Table 4.8: Smith competing overall results

	#	# correct	% correct	# false pos/neg
Raw data	2743	1957	71%	786
LLE only (trial 1)	2743	2211	81%	531
SVDD only (min ℓ^2 error coeffs)	2743	2088	76%	655
SVDD only (mixed ℓ^2-ℓ^1 coeffs)	2743	1497	55%	1245

4.1.4.4 Smith Island Class Maps

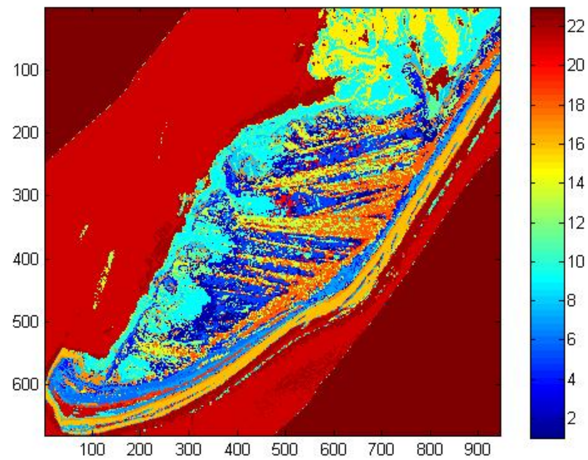

Figure 4.36: Smith trial 1 A class map

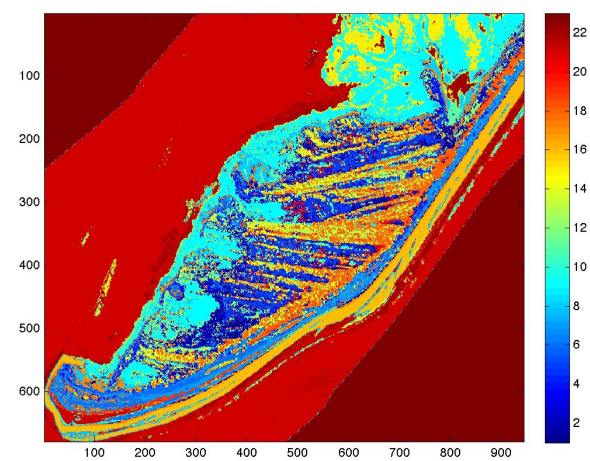

Figure 4.37: Smith trial 1 B class map

4.1.4.5 Smith Island Individual Class Maps

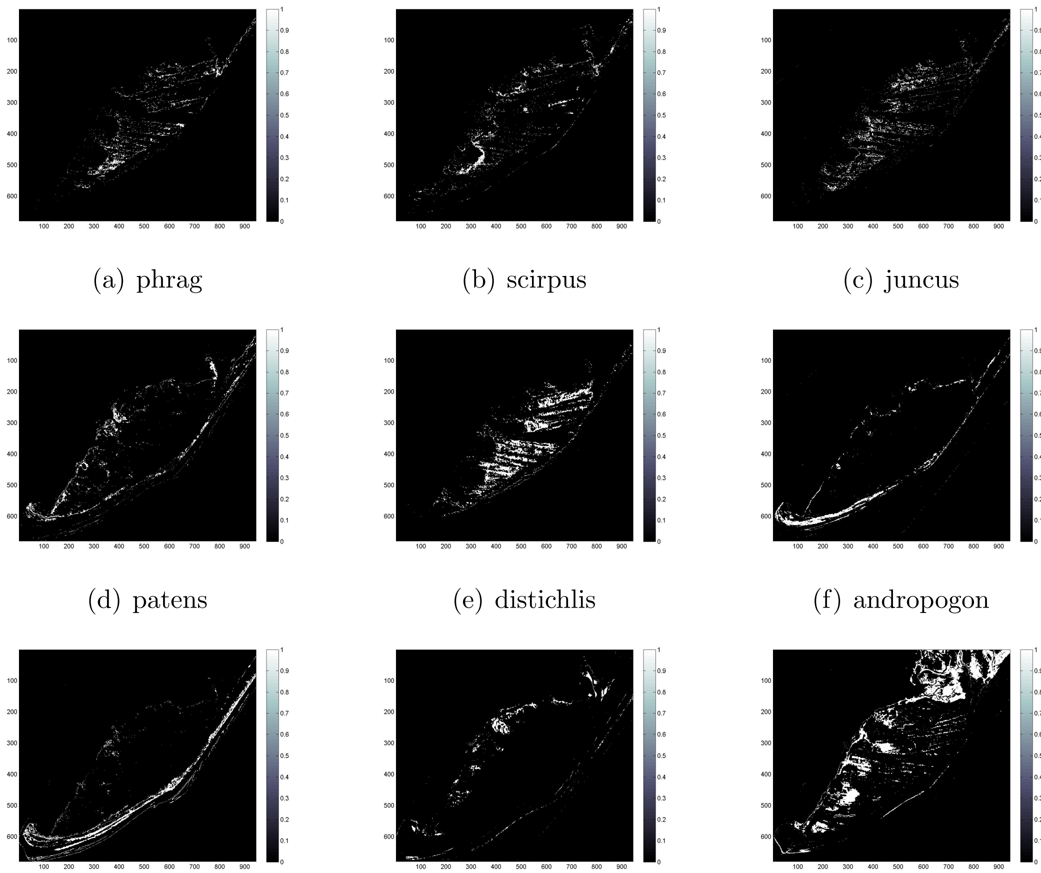

(a) phrag (b) scirpus (c) juncus

(d) patens (e) distichlis (f) andropogon

(g) ammophila (h) mud (i) alterniflora

Figure 4.38: Smith trial 1 A individual class maps 1–9

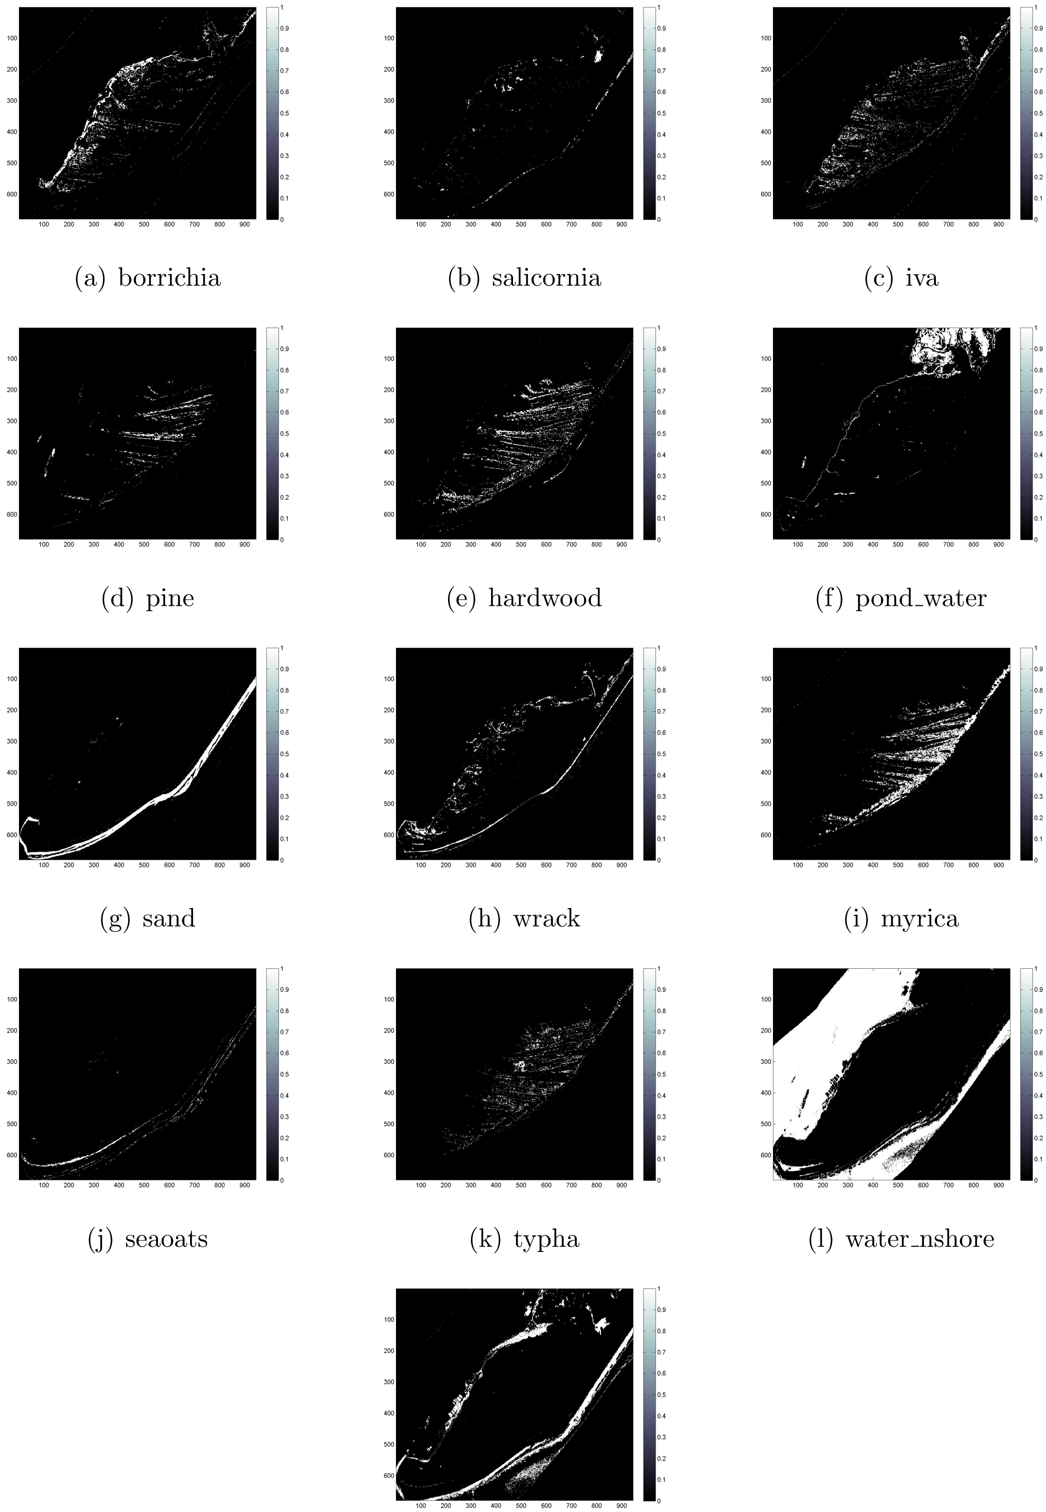

(a) borrichia (b) salicornia (c) iva

(d) pine (e) hardwood (f) pond_water

(g) sand (h) wrack (i) myrica

(j) seaoats (k) typha (l) water_nshore

(m) submerged_nets

Figure 4.39: Smith trial 1 A individual class maps 10–22

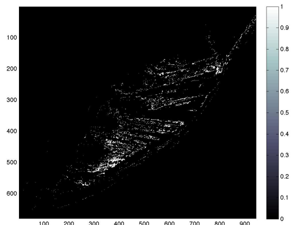

(a) phrag

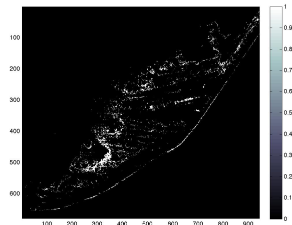

(b) scirpus

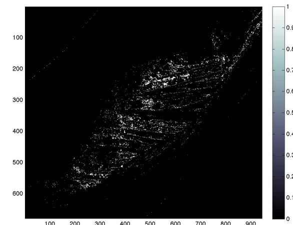

(c) juncus

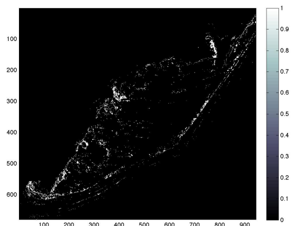

(d) patens

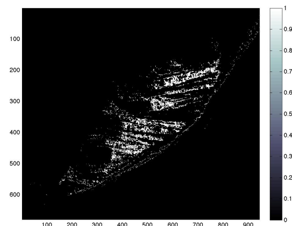

(e) distichlis

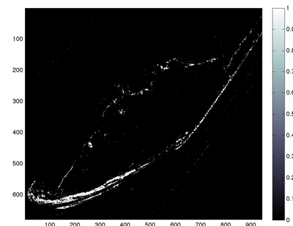

(f) andropogon

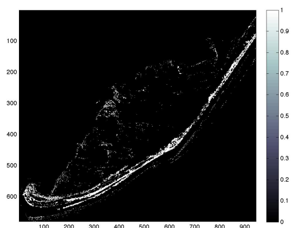

(g) ammophila

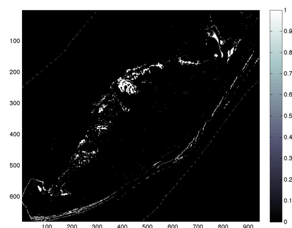

(h) mud

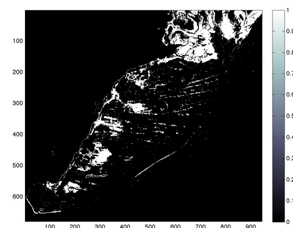

(i) alterniflora

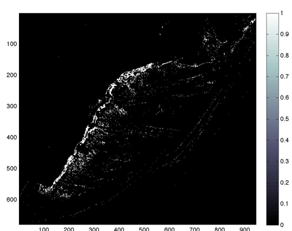

(j) borrichia

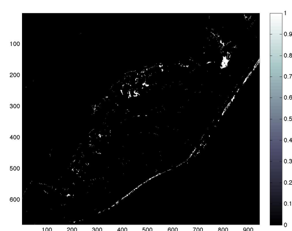

(k) salicornia

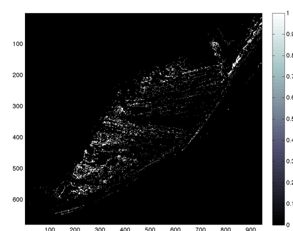

(l) iva

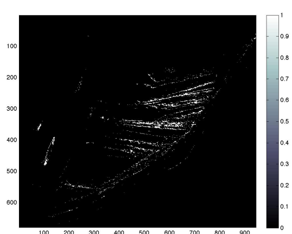

(m) pine

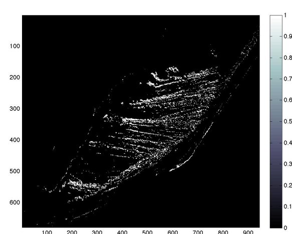

(n) hardwood

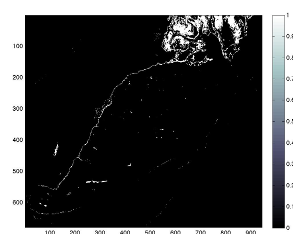

(o) pond_water

Figure 4.40: Smith trial 1 B individual class maps 1–15

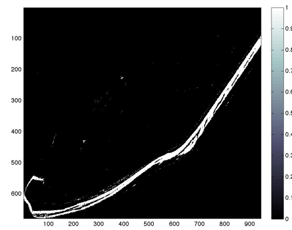

(a) sand

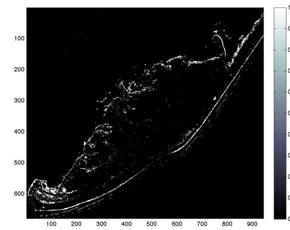

(b) wrack

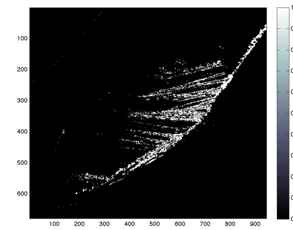

(c) myrica

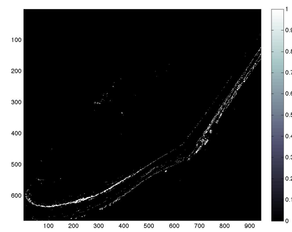

(d) seaoats

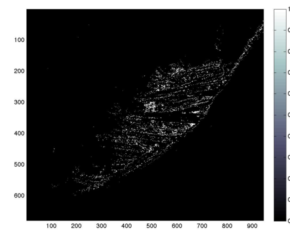

(e) typha

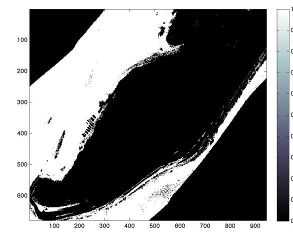

(f) water_nshore

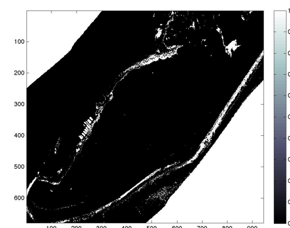

(g) submerged_nets

Figure 4.41: Smith trial 1 B individual class maps 16–22

4.1.4.6 Smith Island Coefficient Maps

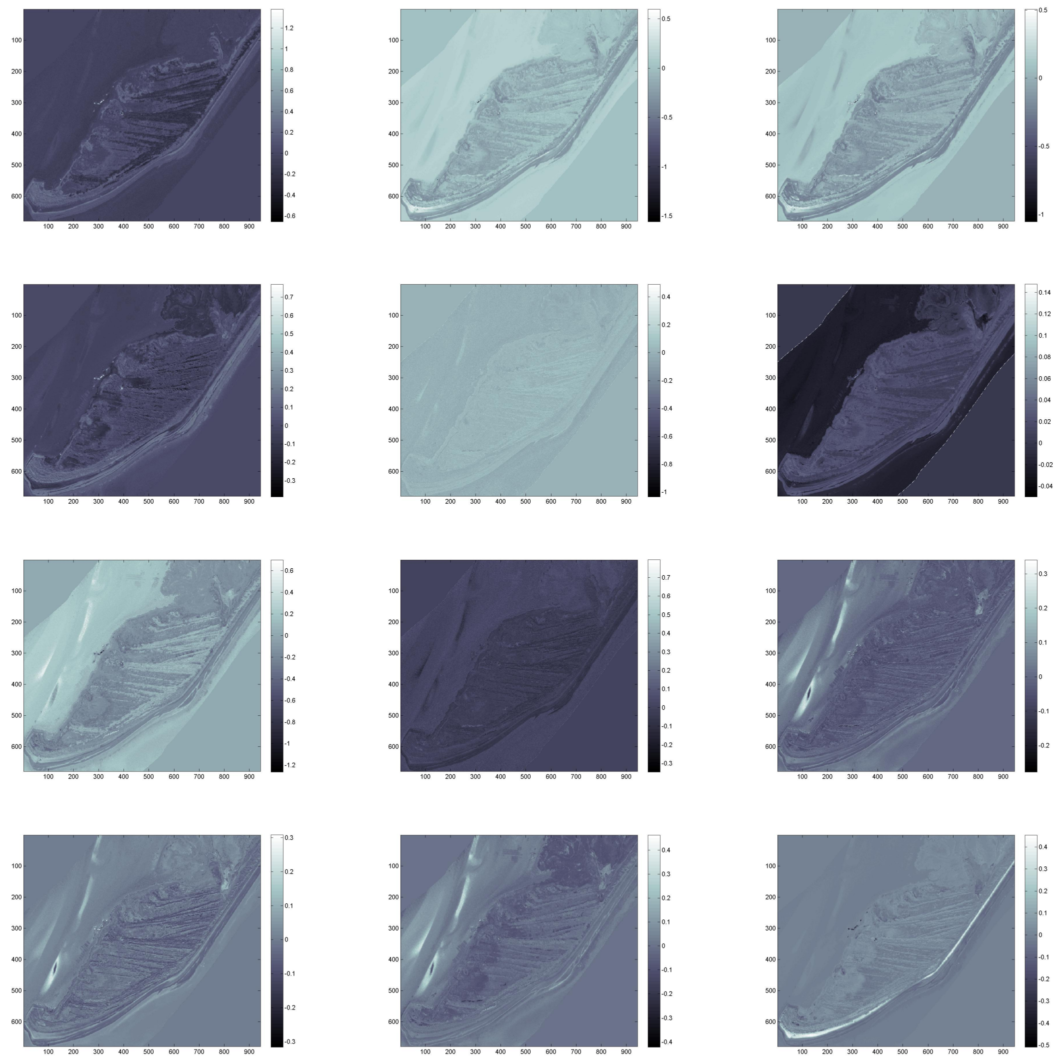

Figure 4.42: Smith trial 1 A canonical coefficients 1–12

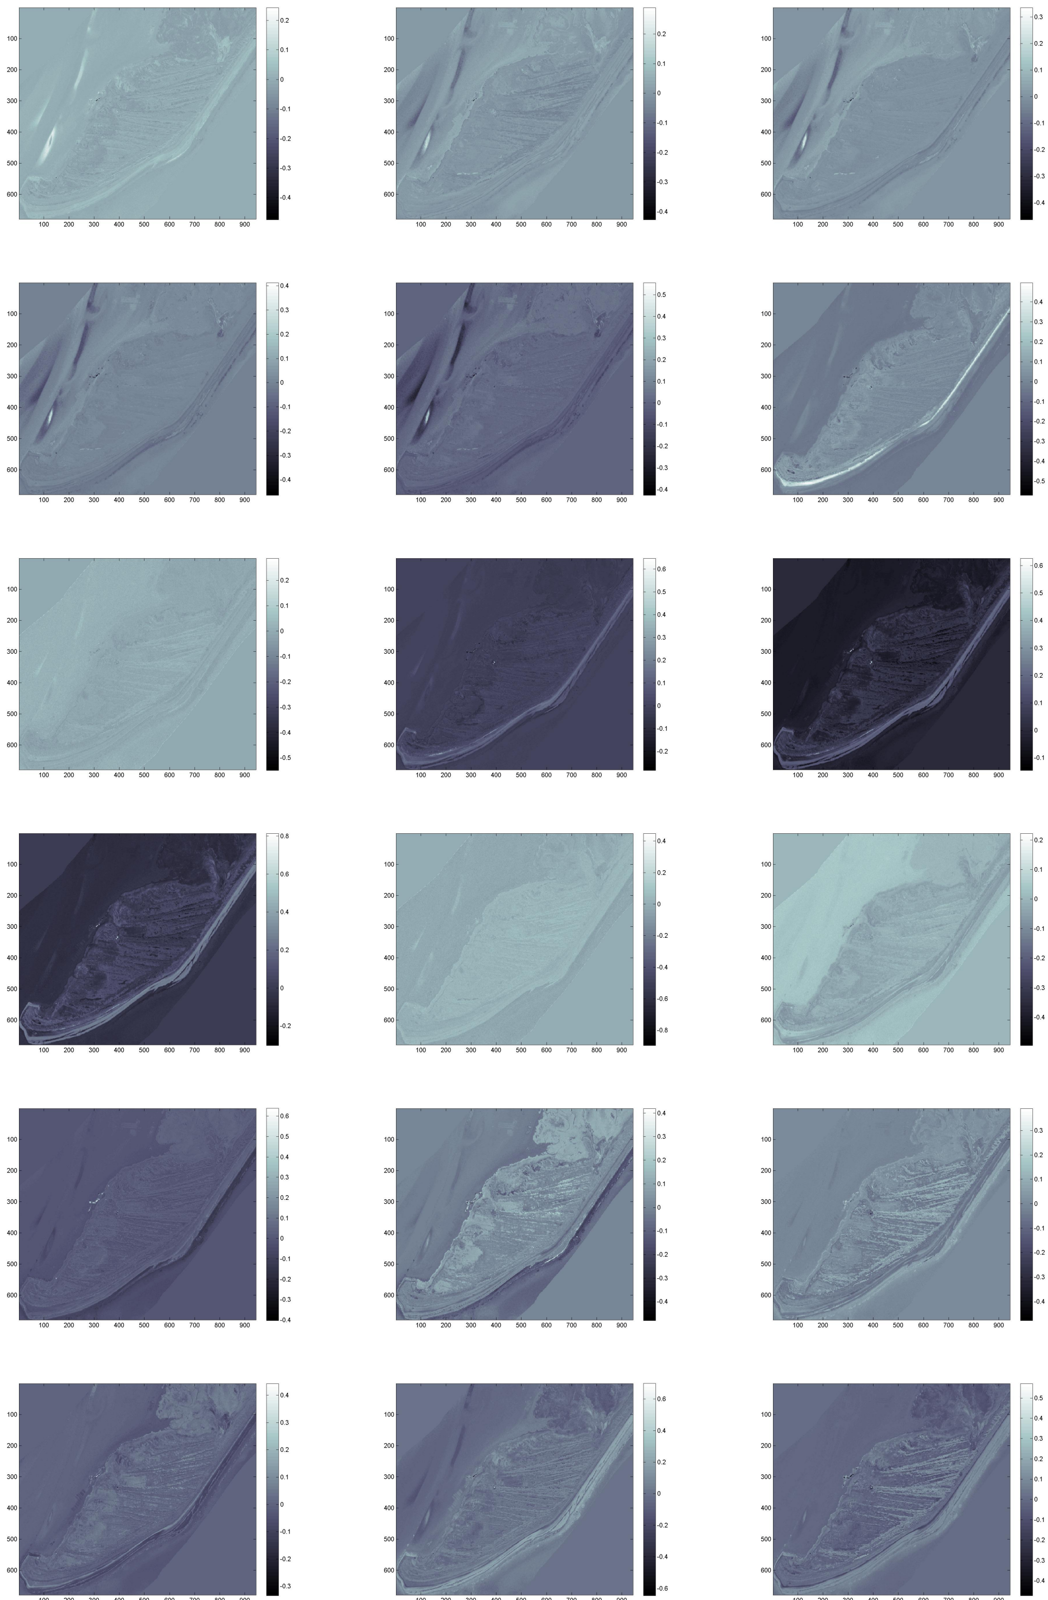

Figure 4.43: Smith trial 1 A canonical coefficients 13–30

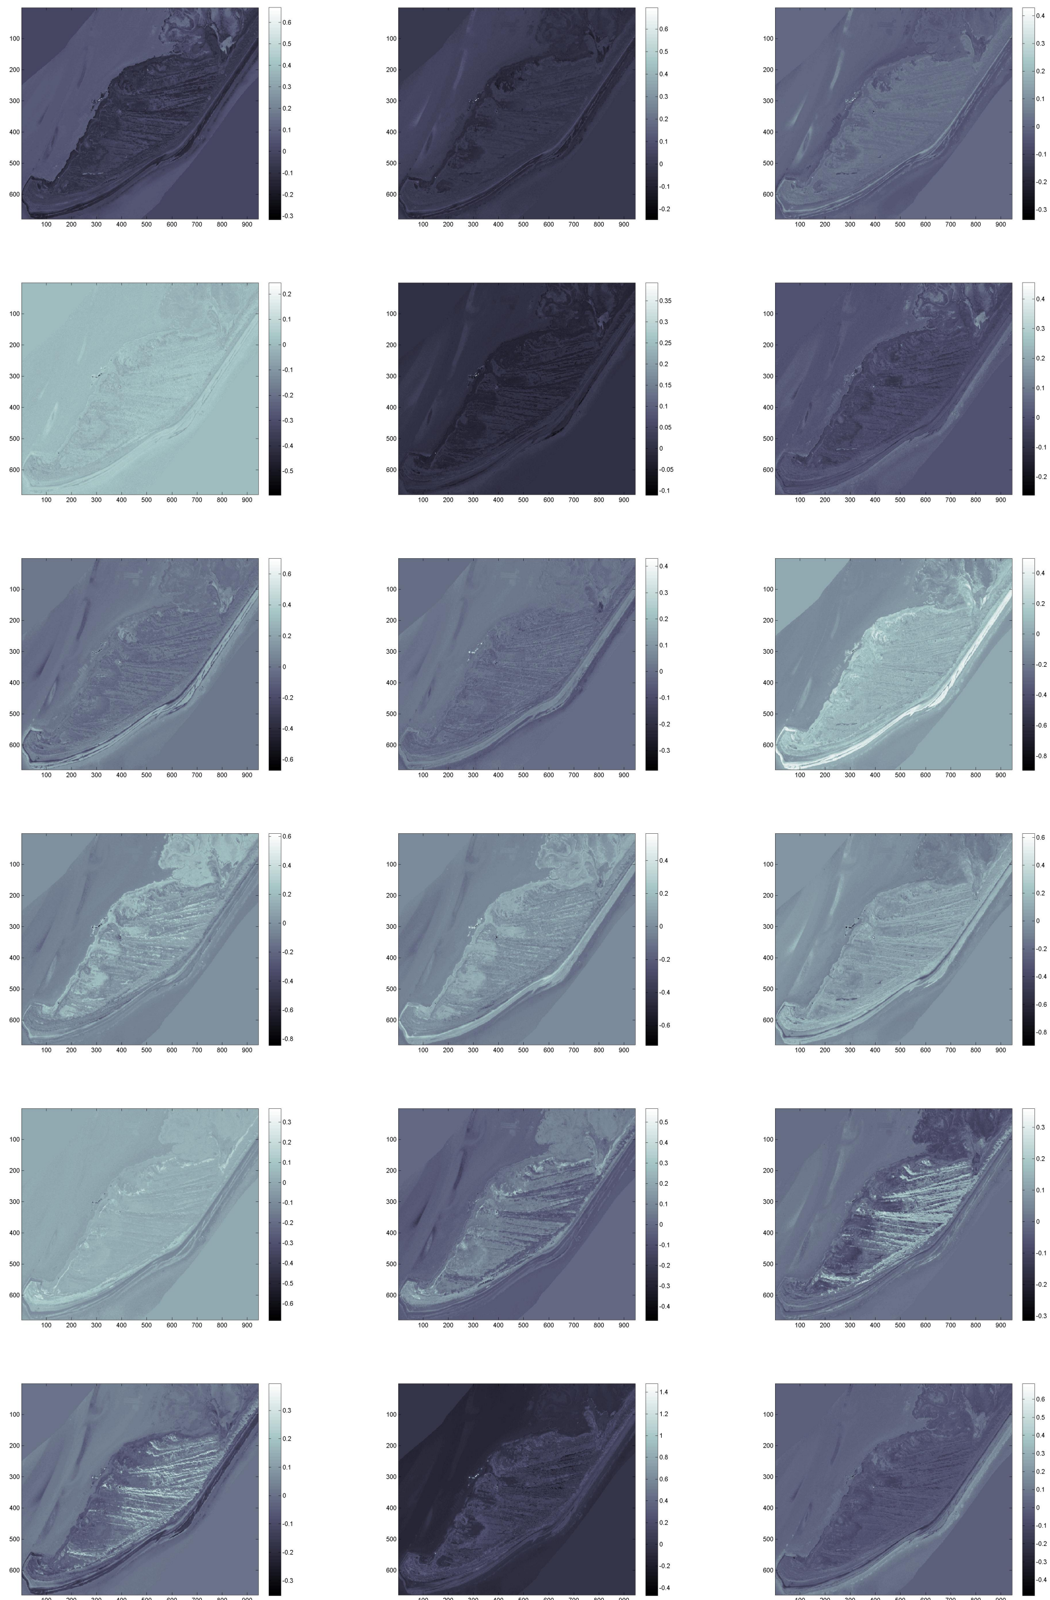

Figure 4.44: Smith trial 1 A canonical coefficients 31–48

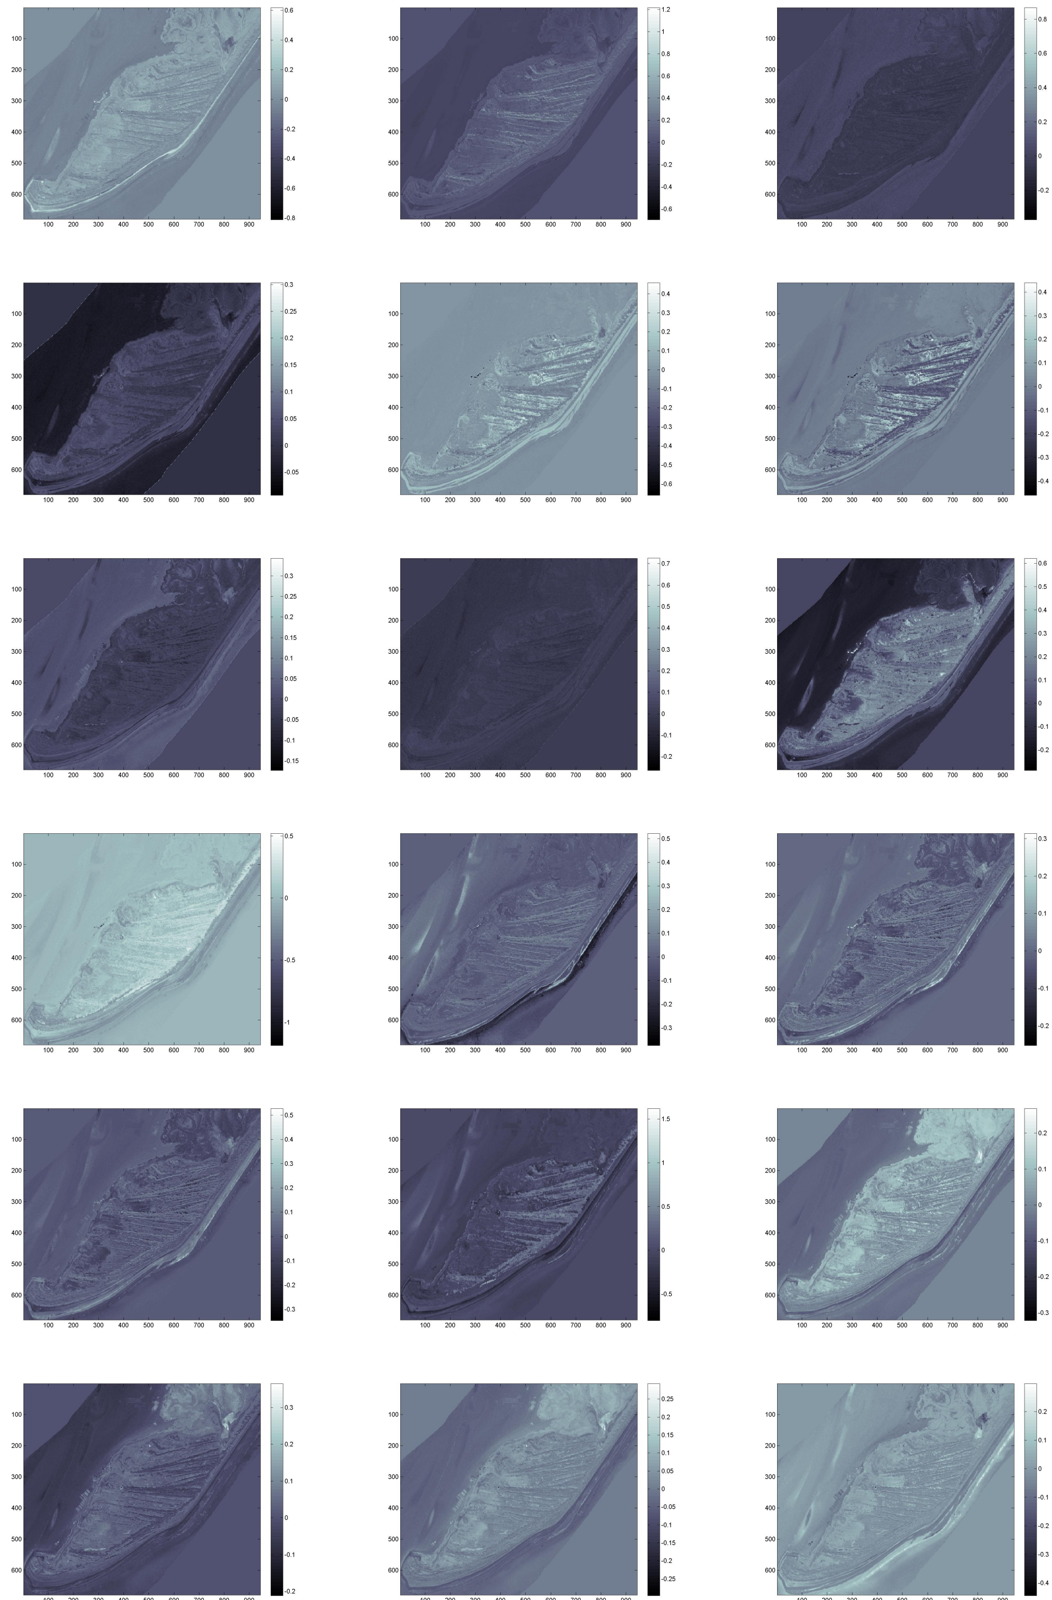

Figure 4.45: Smith trial 1 A canonical coefficients 49–66

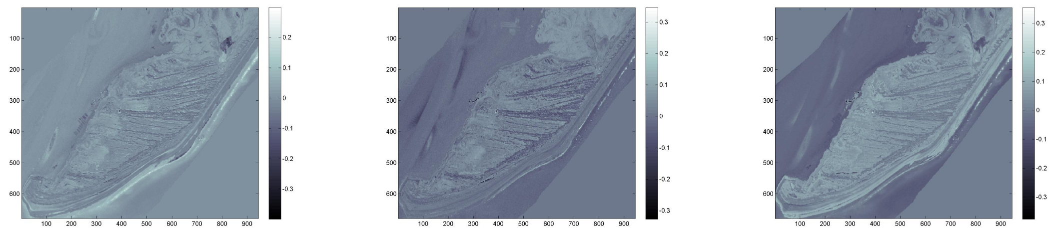

Figure 4.46: Smith trial 1 A canonical coefficients 67-69

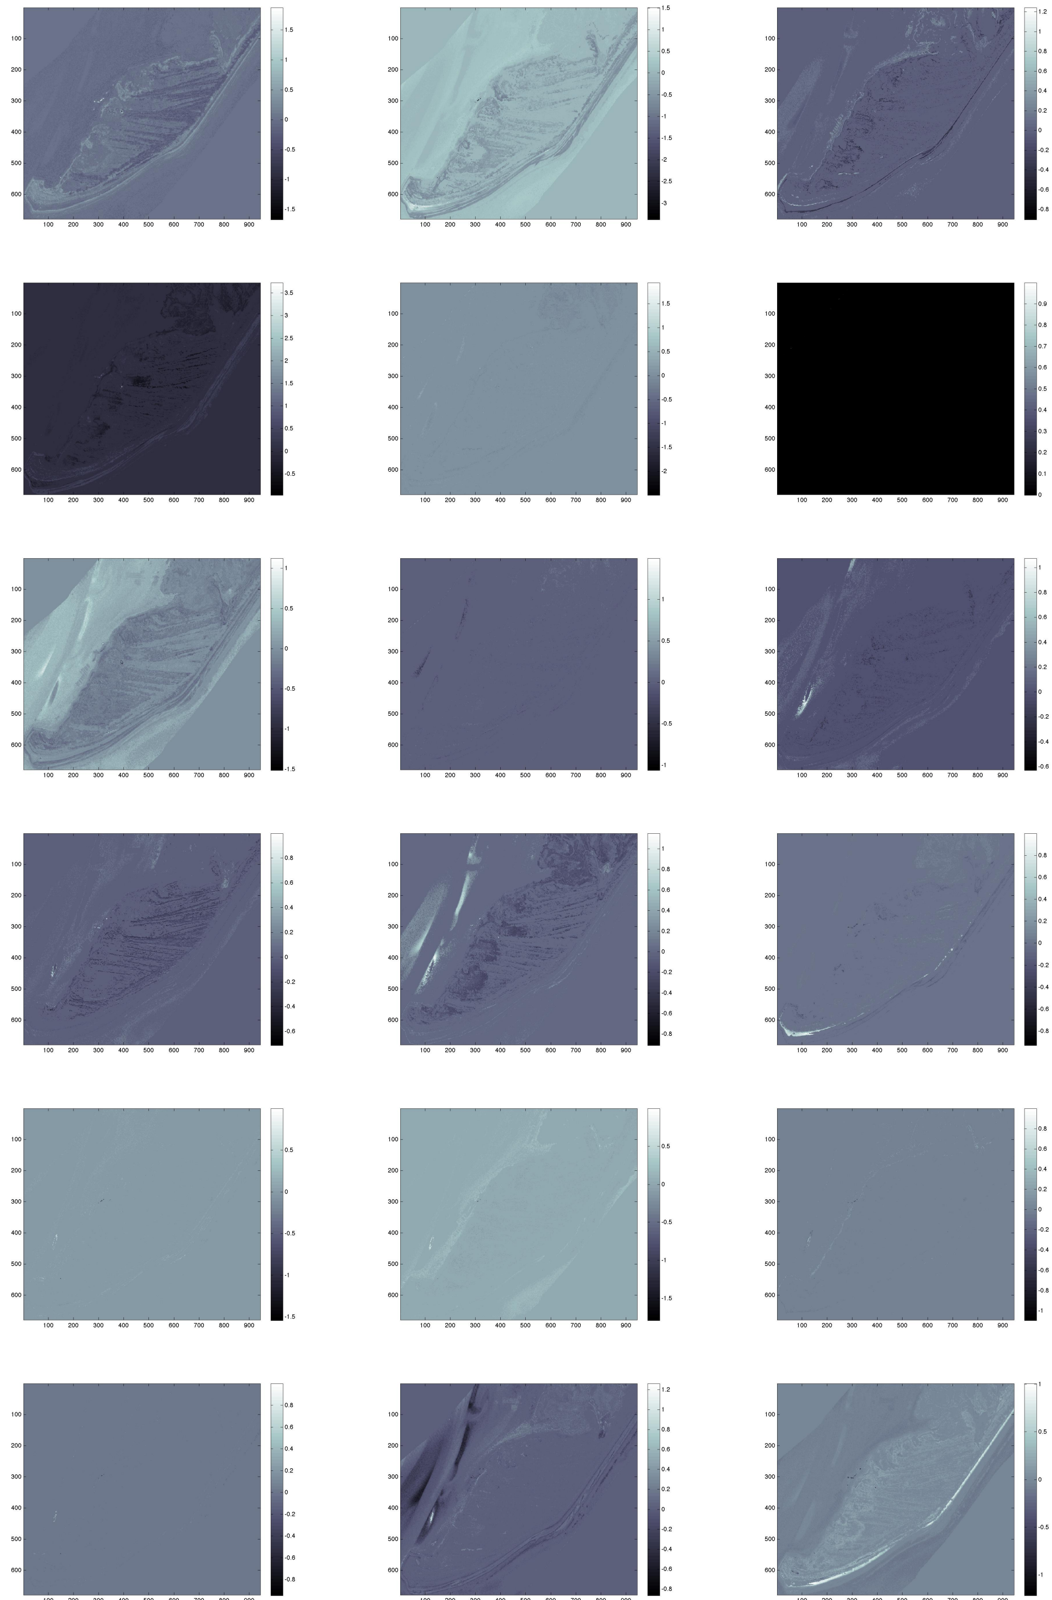

Figure 4.47: Smith trial 1 B sparse coefficients 1–18

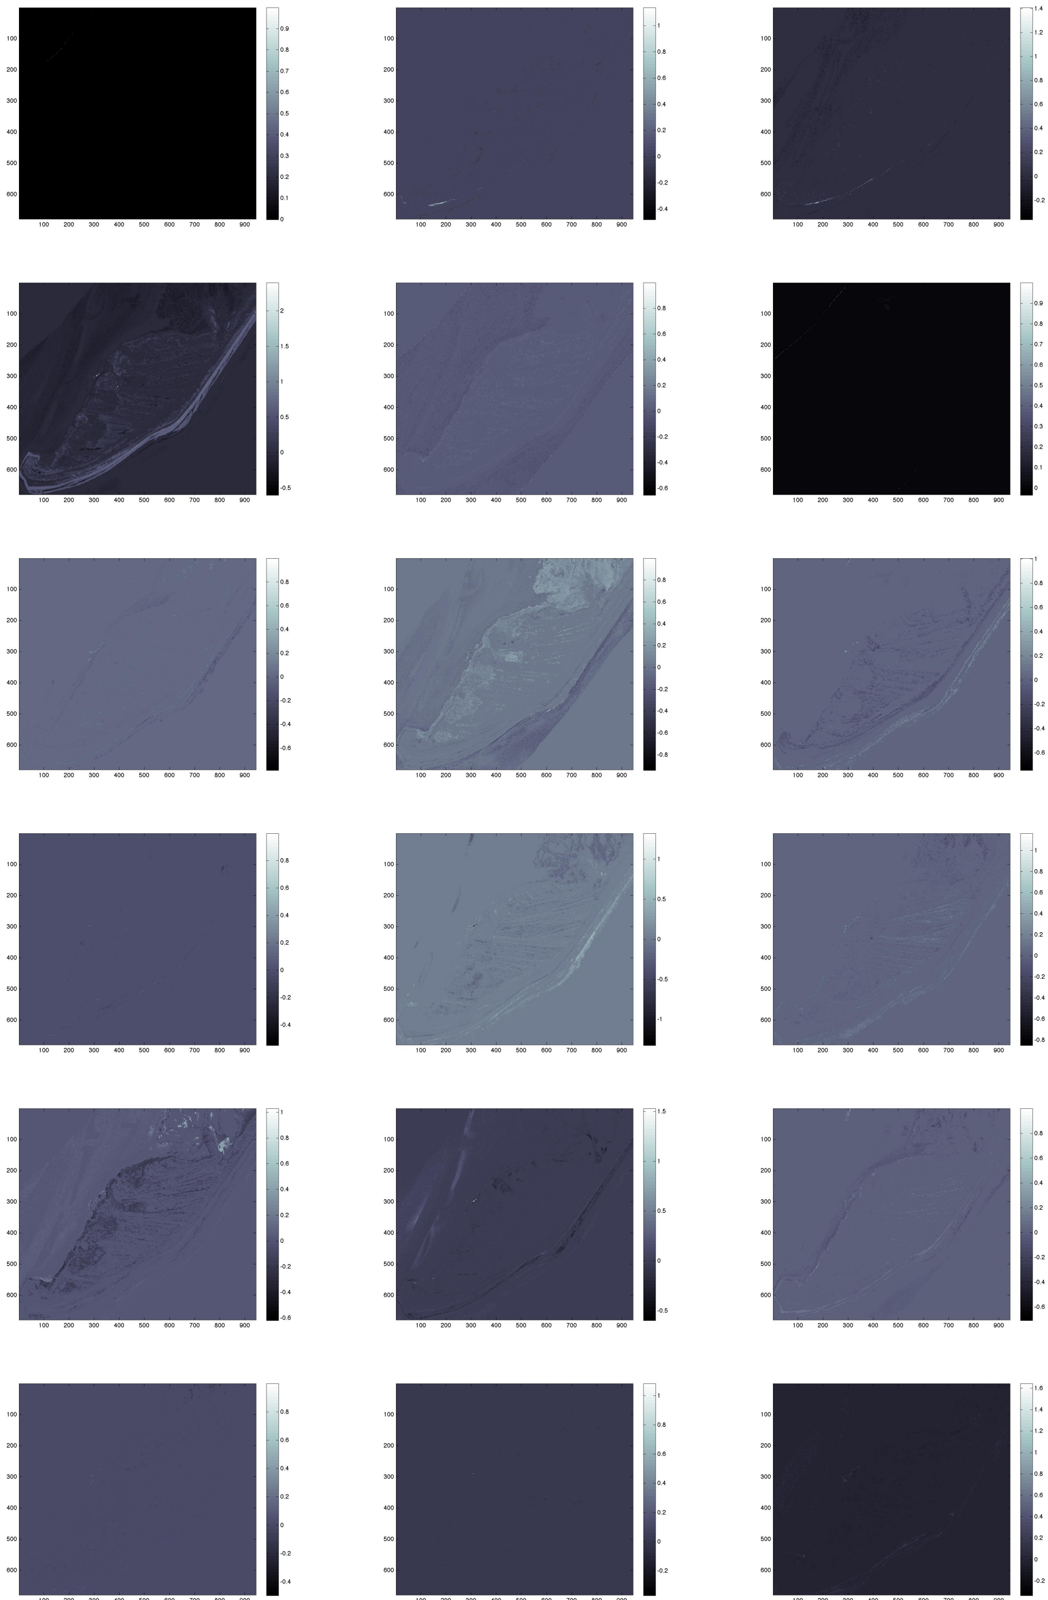

Figure 4.48: Smith trial 1 B sparse coefficients 19–36

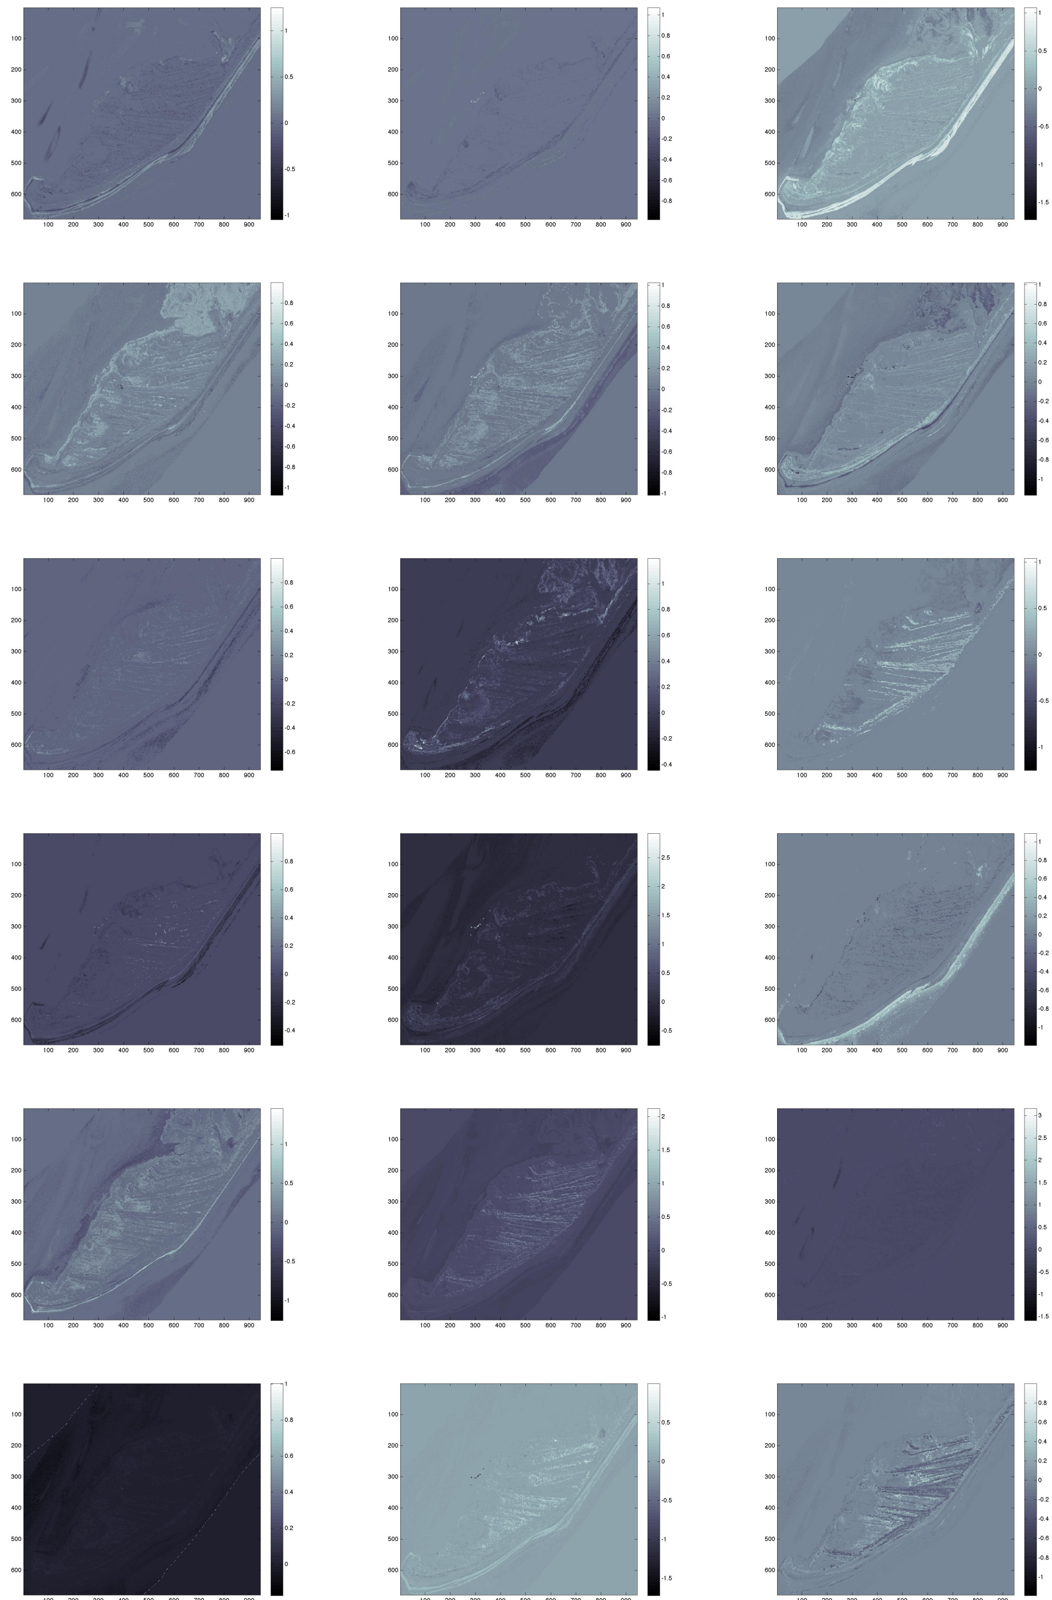

Figure 4.49: Smith trial 1 B sparse coefficients 37–54

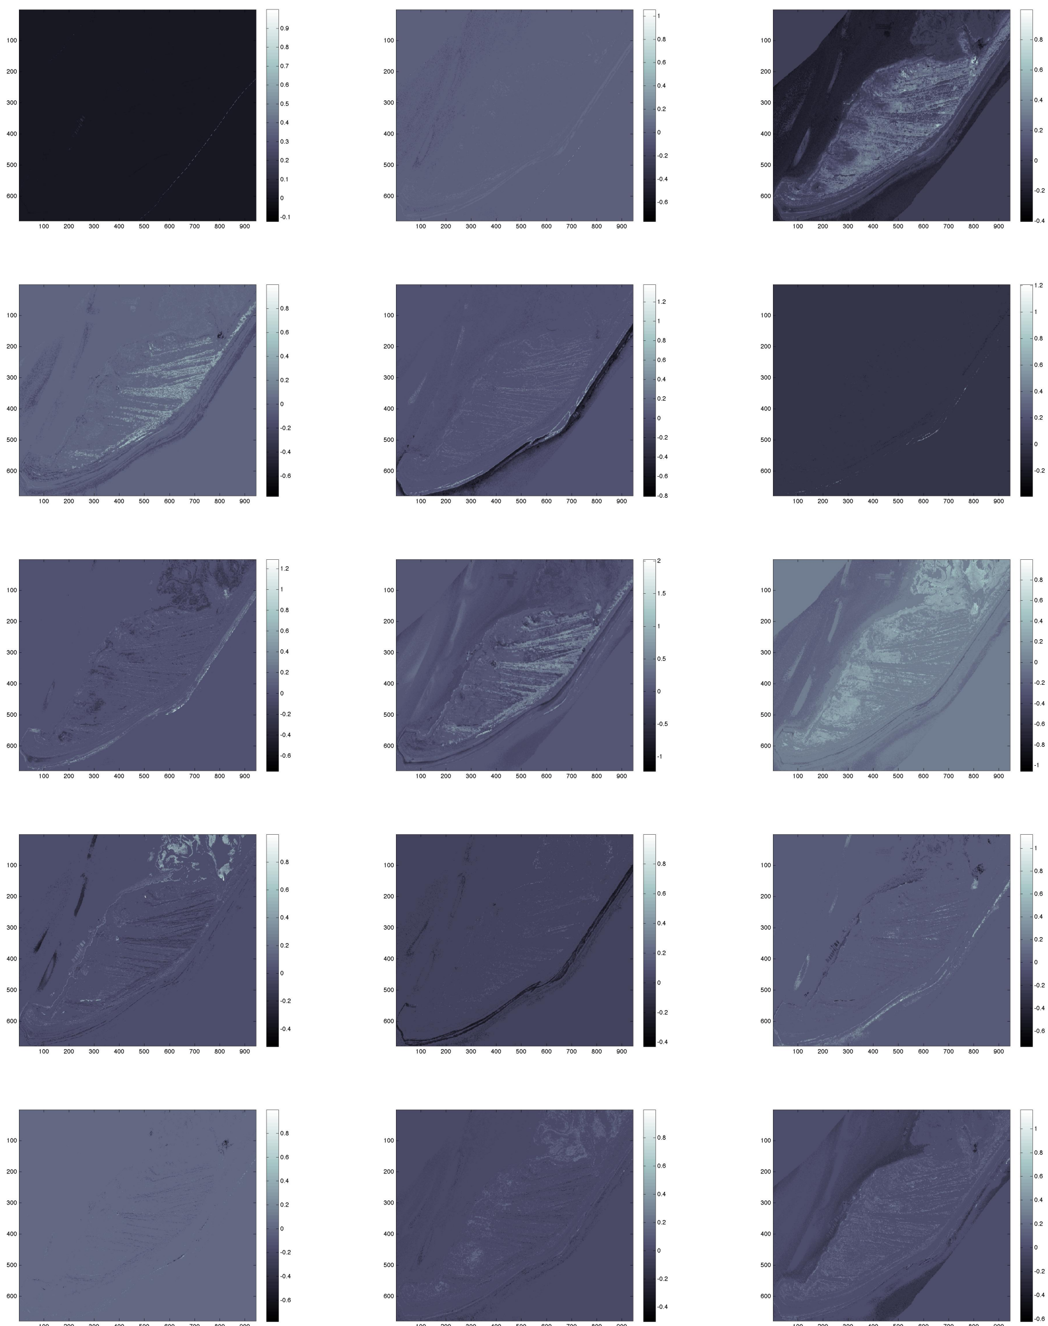

Figure 4.50: Smith trial 1 B sparse coefficients 55–69

4.1.5 Conclusions

The numerical statistics show that, for the most part, our algorithm improves upon the raw data and the SVDD method, while remaining even with LLE, at least on the Urban data set. The one exception comes from the LLE coordinates of Smith Island, which when $d = 43$, LLE alone attains the best classification results. It should be noted, though, that when restricted to the range $10 \leq d \leq 25$, the LLE coordinates on Smith Island attain a maximum classification percentage of 75%, below what our algorithm attains in the same range for d. Similarly, for the Urban data set, larger values of d seem to generate the best classification results for LLE alone, but values of d near the number of classes or below seem to work best when incorporated into our algorithm. Given that both data sets have 22 distinct classes, it would seem plausible that they would lie closer to a manifold with dimension somewhere near 22 as opposed to one with a dimension in the 40's. Yet with LLE alone the spectral angle classifier desires more and more dimensions, as opposed to our algorithm which seems to prefer a more 'appropriate' low dimensional space. Of course this could also be due to the fact that as d increases, the redundancy of the frame lessens, thus reducing its effectiveness.

Another point to make, however, is the argument of storage versus speed. What we have gained in reducing the size of our data set, we have lost in speed. It is rather quick to run SVDD on the original data sets, even for a large number of different parameters. However, even with the speed ups employed in the kernel process, such as landmarking and the out of sample extension, computing a kernel on Urban, and especially Smith, is a time consuming process that can take upwards of 24 hours even on a 8 core computer with 16 gigabytes of RAM. Concerning our algorithm, one must weigh the benefits of reduced storage and increased precision at the cost of time. Of course our algorithm does manage to do more with less, at least compared to LLE, and so perhaps can serve as a compromise between endmember algorithms and kernel methods.

Another point in favor of frames, though, comes from figures 4.2, 4.3, and 4.35. These graphs clearly show a drop off in classification results as one goes from having an over-complete frame to an under-complete endmember set, at least when dealing with the reduced coordinates Y. Perhaps the same is true in the original space X, but given the high dimension that X lies in it is hard to construct a frame with the same redundancy of those constructed for Y.

One final comparison comes from the type of coefficients. In one group we have the minimum ℓ^2-error coefficients (for endmembers) and the canonical coefficients (for frames), and in the other group we have the mixed ℓ^2-ℓ^1 coefficients (for endmembers) and the sparse coefficients (for frames). Despite the added complexity and increased visual appeal of the sparse types of coefficients, it is the simpler ℓ^2 and canonical coefficients that did better in terms of classification. Perhaps for material classification the ℓ^2/canonical coefficients a preferable, but for material identification, one should go with the sparse coefficients.

4.2 Multispectral Retinal Data

The purpose of this experiment is to aid in research concerning age related macular degeneration (AMD), which is one of the leading causes of blindness in the elderly population. One of the indicators of AMD is the presence of irregular lipofuscine deposits, also known as drusin. Using our techniques developed in chapter 3, we present an automated method for the early detection of drusin in retinal imagery.

In order to apply our techniques, we need a high dimensional data set. Through the National Institute of Health, we have obtained a multispectral retinal imagery data set, known to contain drusen. We apply the Laplacian eigenmap kernel to obtain low dimensional coordinates Y. We then construct a maximally separated frame Φ for span(Y), and represent each $y_i \in Y$ in terms of sparse coefficients C. The goal is to have certain frame elements correspond to the drusin. We have one trial on the retinal data to illustrate this process.

4.2.1 Description of the Retinal Data Set

We work on $500 \times 500 \times 20$ patch of the retinal data set; that is $500 * 500 = 250000$ pixels and 20 spectral bands. A color image of the entire data set is displayed in figure 4.51, with the patch that we work on cut out; that patch is magnified in figure 4.52. A sample band of the data set is given in figure 4.53. The only class that we interested in finding is the drusen; there is no ground truth.

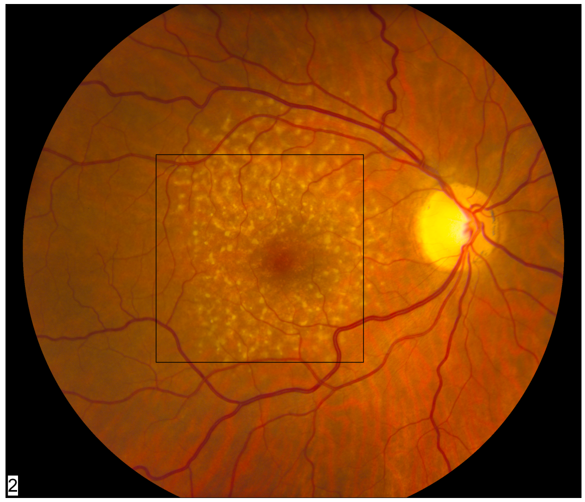

Figure 4.51: Color image of entire retinal data set

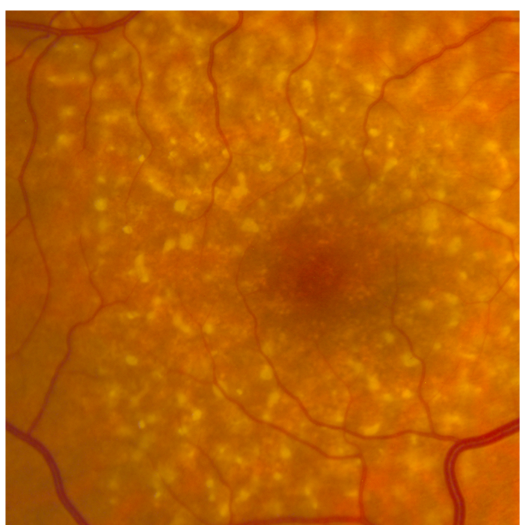

Figure 4.52: Magnified color image patch

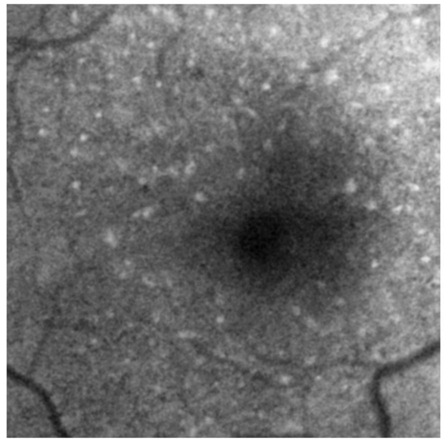

Figure 4.53: Sample band of retinal data set

4.2.2 Retinal Data Trial 1

The results of retinal data trial 1 were obtained with the following settings:

- Data set: X = retinal data
- Kernel: Laplacian eigenmaps
- Number of neighbors: $k = 12$
- Laplacian eigenmaps sigma parameter: $\sigma = 1$
- Frame construction: maximally separated frame
- Number of reduced dimensions: $d = 7$
- Number of frame elements: $s = 15$
- t parameter in (3.5.64): $t = 12$

- ε parameter in (3.5.64): $\varepsilon = .02$

- Type of coefficients: sparse

Figures 4.54, 4.55, and 4.56 show the coefficient maps for each of the frame elements.

4.2.3 Retinal Data Coefficient Maps

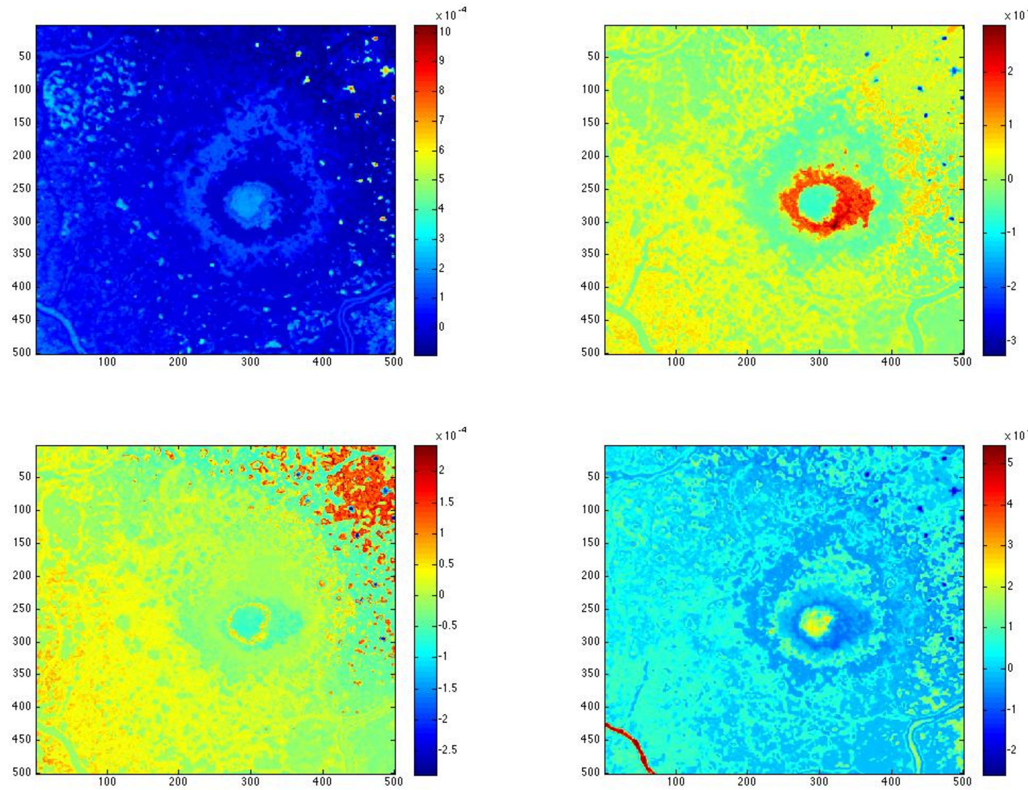

Figure 4.54: Retinal data trial 1 sparse coefficients 1–4

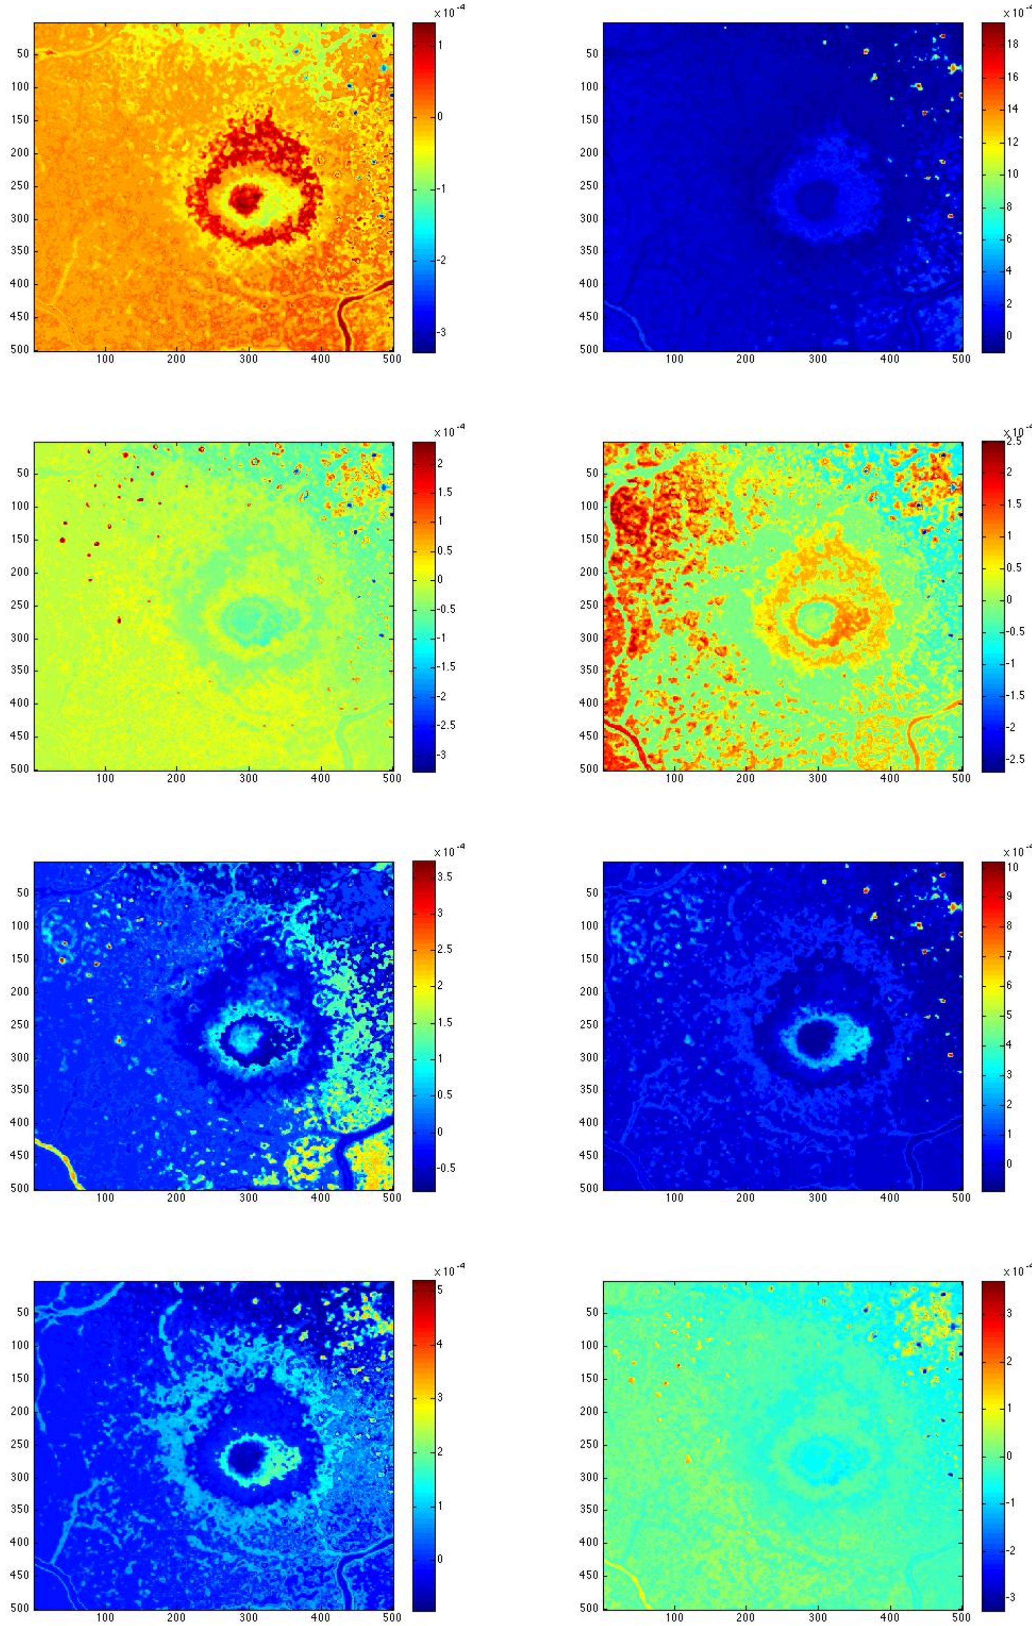

Figure 4.55: Retinal data trial 1 sparse coefficients 5–12

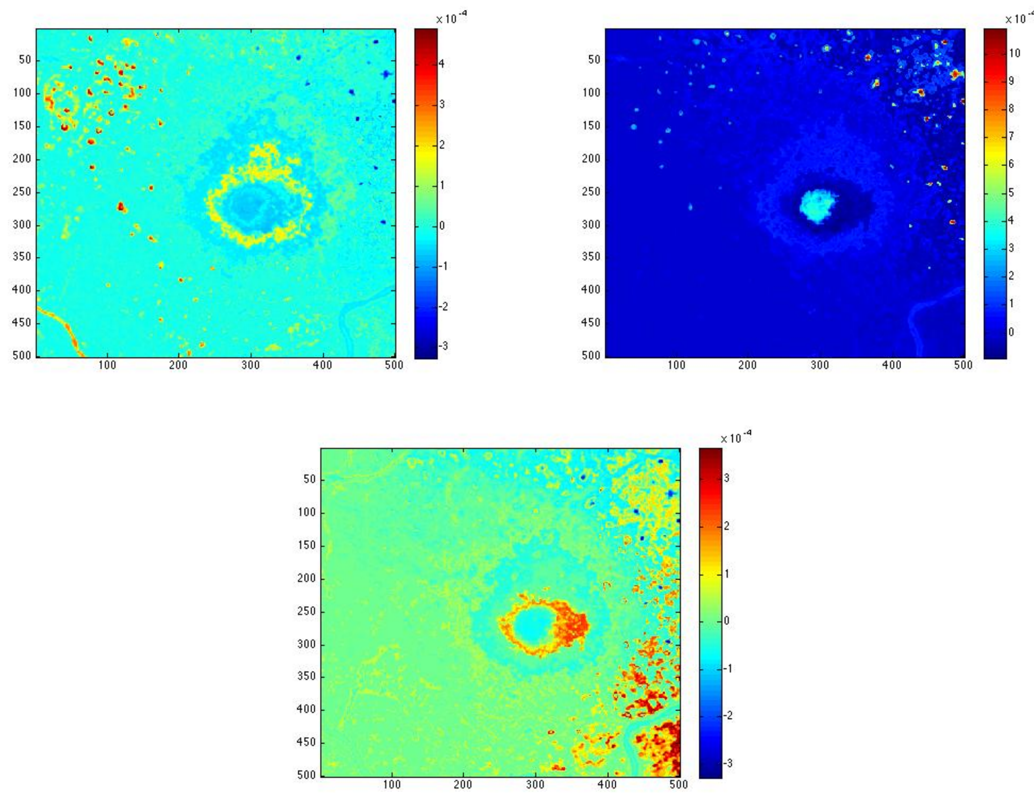

Figure 4.56: Retinal data trial 1 sparse coefficients 13–15

4.2.4 Conclusions

The first thing to note is that we set $t = 12$ in (3.5.64). The idea behind this choice is that the drusin in terms of area in the image are small, and they also have low spectral intensity. Given this knowledge, we allowed 12 of the 15 frame elements to have as large as correlation with the data as they liked, but limited the remaining three to only a 2% correlation ($\varepsilon = .02$). Given the physical nature of the drusin and the fact that they are unique, yet make up an extremely small part of the data, it seemed logical that those three restricted frame elements would gravitate towards the drusin, as opposed to some larger feature. Indeed, of the last three coefficient maps, two of them rather clearly mark drusin (the top two maps in figure 4.56). Furthermore, these two maps seem to mark two separate categories of drusin, the left map marking drusin on the left side of the image, the right map marking drusin on the right side of the image. Perhaps this speaks to a difference between early and later stage drusin, and/or chemical differences. Further investigation is necessary to know for sure.

Bibliography

[1] C. M. Bachmann. Improving the performance of classifiers in high-dimensional remote sensing applications: An adaptive resampling strategy for error-prone exemplars (aresepe). *IEEE Trans. on Geoscience and Remote Sensing*, vol. 41 (9):2101–2112, 2003.

[2] C. M. Bachmann, T. L. Ainsworth, and R. A. Fusina. Exploiting manifold geometry in hyperspectral imagery. *IEEE Trans. on Geoscience and Remote Sensing*, vol. 43 (3):441–454, 2005.

[3] C. M. Bachmann, T. L. Ainsworth, and R. A. Fusina. Improved manifold coordinate representations of large scale hyperspectral imagery. *IEEE Trans. on Geoscience and Remote Sensing*, vol. 44 (10):2786–2803, 2006.

[4] C. M. Bachmann, T. F. Donato, R. A. Fusina, M. H. Bettenhausen, A. L. Russ, J. Burke, G. M. Lamela, W. J. Rhea, B. R. Truitt, and J. H. Porter. A credit assignment approach to fusing classifiers of multi-season hyperspectral imagery. *IEEE Trans. on Geoscience and Remote Sensing*, vol. 41 (11):2488–2499, 2003.

[5] C. M. Bachmann, T. F. Donato, G. M. Lamela, W. J. Rhea, M. H. Bettenhausen, R. A. Fusina, K. Du Bois, J. H. Porter, , and B. R. Truitt. Automatic classification of land-cover on smith island, va using hymap imagery. *IEEE Trans. on Geoscience and Remote Sensing*, vol. 40 (10):2313–2330, 2002.

[6] Amit Banerjee, Philippe Burlina, and Joshua Broadwater. A machine learning approach for finding hyperspectral endmembers. *IEEE International Geoscience and Remote Sensing Symposium*, IGARSS 2007:3817–3820, 2007.

[7] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Advances in Neural Information Processing Systems*, 14:585–591, 2002.

[8] John J. Benedetto and Matthew Fickus. Finite normalized tight frames. *Adv. Comput. Math.*, 18:357–385, 2003.

[9] Yoshua Bengio, Jean-Francois Paiement, and Pascal Vincent. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. *Advances in Neural Information Processing Systems*, 16:177–184, 2004.

[10] Joseph Boardman, Fred Kruse, and Robert Green. Mapping target signatures via partial unmixing of aviris data. *Fifth JPL Airborne Earth Science Workshop, volume 1 of JPL Publication 95-1*, pages 23–26, 1995.

[11] Jeffrey Bowles, Peter Palmadesso, John Antoniades, Mark Baumbeck, and Lee Rickard. Use of filter vectors in hyperspectral data analysis. *Proceedings of SPIE*, 2553:148–157, 1995.

[12] Emmanuel Candes, Justin Romberg, and Terrence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.

[13] Emmanuel Candes and Terrence Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51:4203–4215, 2005.

[14] Peter Casazza and Jelena Kovacevic. Equal-norm tight frames with erasures. *Adv. Comput. Math.*, 18(2–4):387–430, 2003.

[15] Chein-I Chang, Chao-Cheng Wu, Wei min Liu, and Yen-Chieh Ouyang. A new growing method for simplex-based endmember extraction algorithm. *IEEE Transactions on Geoscience and Remote Sensing*, pages 2804–2819, October 2006.

[16] Ole Christensen. *An Introduction to Frames and Riesz Bases*. Applied and Numerical Harmonic Analysis. Birkhauser, Boston, 2003.

[17] David Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306, April 2006.

[18] David Donoho and Carrie Grimes. Hessian eigenmaps: New locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 102(21):7426–7431, 2005.

[19] David Dummit and Richard Foote. *Abstract Algebra*. John Wiley and Sons, Inc., 3rd edition, 2004.

[20] Yonina C. Eldar and Helmut Bolcskei. Geometrically uniform frames. *IEEE Trans. Inform. Theory*, 49(4):993–1006, 2003.

[21] David Gillis, Jeffrey Bowles, and Michael Winter. Using endmembers as a coordinate system in hyperspectral imagery. *Proceedings of SPIE*, 4816:346–354, 2002.

[22] Vivek Goyal, Jelena Kovacevic, and Jonathan Kelner. Quantized frame expansions with erasures. *Journ. of Appl. and Comput. Harmonic Analysis*, 10(3):203–233, 2001.

[23] Ian Jolliffe. *Principal Component Analysis*. Springer, 2nd edition, 2002.

[24] Jelena Kovacevic and Amina Chebira. Life beyond bases: The advent of frames (part I). *IEEE SP Mag.*, 24(4):86–104, July 2007. Feature article.

[25] Jelena Kovacevic and Amina Chebira. Life beyond bases: The advent of frames (part II). *IEEE SP Mag.*, 24(5):115–125, September 2007. Feature article.

[26] John Mercer. Functions of positive and negative type and their connection with the theory of integral equations. *Philosophical Transactions of the Royal Society of London*, 1909.

[27] Robert Reams and Shayne Waldron. Isometric tight frames. *Electron. J. Linear Algebra*, 9:122–128, 2002.

[28] Sam Roweis and Lawrence Saul. Nonlinear dimension reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

[29] John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis.* Cambridge University Press, Cambridge, UK, 2004.

[30] Peg Shippert. Why use hyperspectral imagery. *Photogrammetric Engineering and Remote Sensing*, pages 377–380, April 2004.

[31] David Tax and Robert Duin. Support vector data description. *Machine Learning*, 54:45–66, 2004.

[32] Joshua Tenenbaum, Vin de Silva, and John Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[33] Richard Vale and Shayne Waldron. Tight frames and their symmetries. *Const. Approx.*, 21(1):83–112, 2005.

[34] Shayne Waldron and Nick Hay. On computing all harmonic frames of n vectors in $\mathbb{C}^d$. *Appl. Comput. Harmon. Anal.*, 21:168–181, 2006.

[35] Michael Winter. N-FINDR: An algorithm for fast autonomous spectral endmember determination in hyperspectral data. *Proceedings of SPIE*, 3753:266–275, 1999.

[36] Georg Zimmermann. Normalized tight frames in finite dimensions. In W. Haussmann, K. Jetter, and M. Reimer, editors, *Recent Progress in Multivariate Approximation*, pages 249–252. Internat. Ser. Numer. Math., 2001.